Published in final edited form as:
Plasmid. 2012 March ; 67(2): 167–173. doi:10.1016/j.plasmid.2012.01.004.

Evaluating quantitative methods for measuring plasmid copy numbers in single cells

Shay Tal and **Johan Paulsson***
Department of Systems Biology, Harvard Medical School, 200 Longwood Ave, Boston, MA 02115, USA

Abstract

The life of plasmids is a constant battle against fluctuations: failing to correct copy number fluctuations can increase the plasmid loss rate by many orders of magnitude, as can a failure to more evenly divide the copies between daughters at cell division. Plasmids are therefore long-standing model systems for stochastic processes in cells, much thanks to the efforts of Kurt Nordström to whose memory this issue is dedicated. Here we analyze a range of experimental methods for measuring plasmid copy numbers in single cells, focusing on challenges, trade-offs and necessary experimental controls. In particular we analyze published and unpublished strategies to infer copy numbers from expression of plasmid-encoded reporters, direct labeling of plasmids with fluorescent probes or DNA binding proteins fused to fluorescent reporters, PCR based methods applied to single cell lysates, and plasmid-specific replication arrest. We conclude that no method currently exists to measure plasmid copy numbers in single cells, and that most methods instead inadvertently measure various types of experimental noise. We also discuss how accurate methods can be developed.

Keywords

Plasmid fluctuations; single cell measurements; copy number distributions

1. Introduction

The control of plasmid replication and partitioning have been closely studied in the last few decades, producing detailed molecular interaction maps and quantitative dynamic descriptions for a wide range of plasmids. Many conclusions were extracted from cleverly designed bulk experiments, inferring mechanistic aspects of control despite only observing the average behavior across a population. However, several important questions are impossible to address at the level of averages. For example, the presumed evolutionary driving force behind plasmid replication control is to correct random fluctuations in copy numbers and thereby reduce the plasmid loss rates – sometimes by several orders of magnitude (Nordström, Molin, & Light, 1984). Fluctuations can for example be caused by random segregation, inherently stochastic expression of control molecules, or variations in the concentrations of chaperones, polymerases, ribosomes that all indirectly affect the

*Corresponding author: Address: Department of Systems Biology, Harvard Medical School, 200 Longwood Ave, Boston, MA 02115, USA, johan_paulsson@hms.harvard.edu.

initiation of replication. The optimal design of self-control then depends on the magnitude and time-scales of these various sources of fluctuations. For example, if the greatest threat to maintaining narrow plasmid distributions comes from slowly meandering changes in the intracellular environment, self-correcting feedback is not so dependent on rapid turnover of control molecules but very dependent on highly cooperative self-inhibition mechanisms (Lestas, Vinnicombe, & Paulsson, 2010). If plasmid fluctuations come from stochastic expression of control molecules, highly cooperative self-inhibition mechanisms by contrast destabilize control and create wider distributions. To understand the logic of the control circuits we must thus consider single cell dynamics – accounting for the control of deviations using a framework that exposes those deviations. The same is true for partitioning mechanisms at cell division that produce more even segregation and thereby reduce the risk that all copies by chance end up in the same daughter: without measuring fluctuations in single cells it is impossible to know what challenges partitioning mechanisms face, how well they perform, or how various mechanistic aspects improve performance. Such analyses currently draw on the rate of plasmid loss from plasmid-containing cells, but loss rates are exceedingly difficult to interpret because they depend sensitively on average copy numbers, the width of distributions, and the partitioning statistics.

The last decade has seen an explosion of interest in stochastic processes in single cells, mainly thanks to the availability of new experimental techniques. Many of the most commonly used models and concepts currently in circulation were in fact first introduced in the context of plasmid biology. Despite this broad interest in stochastic effects and the specific interest in plasmids, there are currently no methods to accurately count plasmid copies in single cells, let alone correlating fluctuations in plasmid copy numbers with plasmid expressed control molecules. Here we describe our long-term efforts to determine plasmid copy number distributions. This work traces back to Kurt Nordström – to whom this issue is dedicated – in many ways. To our knowledge Kurt was the first to model the connection between plasmid heterogeneity and the underlying stochastic mechanisms, including negative feedback loops and stochastic partitioning at cell division (Nordström, Molin, & Light, 1984). Kurt also recognized early on that many quantitative plasmid experiments, for example to determine the mode of partitioning at cell division, would be confounded by plasmid distribution effects. One of us (J.P.) can also thank Kurt personally for this research direction, which began with Kurt's lecture on plasmid copy number control in a course on stochastic processes given by Måns Ehrenberg in 1993. My first scientific meeting with Kurt in turn focused on the role of negative experimental results, where I argued against them on principle and Kurt patiently insisted that, from a practical point of view, a careful negative result can be almost as useful as a positive result. Here we have taken his words to the extreme and summarize a decade of our own negative experimental results (read failures) to measure reliably plasmid copy numbers in individual cells.

Failing to establish seemingly reasonable experimental protocols can reflect an inability to do the experiments correctly, an ability to include the appropriate controls, or a lethal combination of the two. Distinguishing between these problems is difficult, so rather than presenting all details we focus on the principles and trade-offs we positively identified in this process. This paper is thus not intended as a review but rather as an analysis motivated by a large number of unpublished experiments. First we discuss previously published approaches to measure plasmid distributions, all of which explicitly acknowledged that measurement errors dominated over the natural fluctuations, and show how challenging such experiments are, regardless of how they are implemented. We then discuss other approaches that have not been published but that appear appropriate for counting plasmids in cells.

2. Methods for measuring plasmid copy numbers in single cells

In this section we analyze potential methods for measuring plasmid copy numbers distributions, including inferring copy numbers from plasmid expressed reporter proteins, tagging plasmids in cells and running qPCR or digital PCR for single cell lysates. We discuss the advantages and disadvantages of each method, the challenges that must be overcome to make them quantitative, and the controls required to demonstrate that a method works.

2.1 Monitoring plasmid expressed protein in single cells

Published methods—To our knowledge the only published attempts to estimate the number of plasmid copies per cell used expression of a plasmid-encoded fluorescent protein (FP) and measurement of the total fluorescence intensity per cell. Cells with twice the copy number should then have twice the fluorescence, and the shape of the protein distribution was argued to reflect the shape of the underlying plasmid distribution.

Anders Løbner-Olesen published the first attempt to use this approach to measure plasmid copy numbers distributions more than a decade ago (Løbner-Olesen, 1999). He recognized that for stable proteins, the current protein level reflects the whole history of changes in gene dosage, with on average half of the proteins made one generation ago. As plasmids fluctuate over time, the levels of the encoded proteins would thus smooth out the upstream plasmid heterogeneity and make it seem smaller than it actually is. He therefore induced FP expression for 25% of a cell cycle, harvested the cells, blocked further expression and waited for the fluorophores to mature before measuring the fluorescence per cell in a flow cytometer. Since the plasmid copy number changes in discrete steps, a well working method would produce equidistant peaks in the intensity distribution corresponding to integer plasmid copy numbers, at least for low copy plasmids such as F or R1. Such peaks were indeed observed for the chromosomally expressed FP, but not for plasmids. In fact, all plasmids in this study – with widely different average copy numbers, partitioning systems, and replication control – displayed very similar and smooth fluorescence distributions. The problem is partly that the approach measures the total superimposed variation from numerous different molecular sources, some related to plasmid partitioning and replication, others to stochastic gene expression or protein partitioning at cell division, and yet others from considering unsynchronized cells. The study also demonstrated a linear relationship between the average fluorescence and the induction period, but not a strict proportionality: there was a high level of fluorescence even before induction, suggesting leaky promoters. In summary the study concluded that the method at best could detect two-fold differences in plasmid copy numbers at the single cell level, and it was primarily used for demonstrating the great increase in fluctuations for replicators with impaired copy number control.

A decade later, Wong *et al* (Wong Ng, Chatenay, Robert, & Piorer, 2010) tried a similar approach using a different type of flow cytometry, which is more suitable for bacterial work, but induced FP production for almost two cell cycles. They similarly failed to see discrete peaks in the intensity histogram and instead attempted to indirectly infer the variance of the underlying plasmid distribution by comparing it to the variance observed for chromosomally expressed FPs. The problem is that the variation from each noise source is amplified or dampened depending on the exact dynamics of the system, which would distort the underlying distribution in different ways because plasmids and chromsomes control their own replication in different ways. Here we summarize what we believe are the key trade-offs and challenges that any such method must overcome. We focus on FPs for simplicity, but the arguments apply also to other reporter systems.

Expression reporters can both exaggerate and underestimate the underlying fluctuations—Fluctuating levels of ribosomes and RNA polymerases can affect the expression of all genes in a cell, fluctuations in specific transcription factors or RNases affect groups of genes, and the probabilistic nature of individual synthesis and degradation events of mRNAs and proteins affect each gene separately. The *average rate* of protein production may thus accurately reflect the *average* plasmid copy number, but the heterogeneity in expression rates in individual cells also reflects spontaneous noise in transcription, changes in gene activity, transcription factors, polymerases, RNases, and ribosomes. In fact, just as gene expression has been used as a tool in studies focusing on plasmid copy numbers, plasmids are often used as tools to study stochastic gene expression (Becskei & Serrano, 2000). Such assays may thus greatly overestimate plasmid fluctuations by mistaking gene expression noise for plasmid noise.

The slow dynamics of FP levels also distort the underlying fluctuations. Imagine that at some time point there is a perfect correlation between plasmid copy numbers and the levels of a plasmid-expressed protein in single cells. When plasmids replicate, the *rates* of transcription and translation may soon increase by the same factor, but the actual *level* of the protein will lag behind, and only asymptotically adjust to the new plasmid level with a time constant set by the total rate of protein degradation and dilution. For stable proteins it would take several generations before the protein has adjusted to its new level. However, on that time scale plasmids replicate several times. Before protein levels reach their new quasi steady state, plasmids will thus have changed again. As a consequence, even if there were no other sources of heterogeneity, measuring the total fluorescence intensity of continuously expressed stable fluorescent protein will not accurately reflect fluctuations in the plasmid copy number, but rather some average over the history of the last few cell generations.

Then at cell division, the plasmids and the FPs are partitioned between the two daughters, with statistical errors that depend on abundances and partitioning mechanisms, but without any expected correlation between plasmids and proteins. For example, if the FPs are present in high numbers they are expected to partition more or less symmetrically between the two daughter cells, while plasmids can display substantial partitioning errors depending on clustering, partitioning systems and copy number. Again the protein level will adjust towards the post-division plasmid level, but too slowly to track the plasmid level without substantial errors. The FP levels are thus in a constant chase after the plasmid, but always lagging behind, which means that, overall, FPs can greatly underestimate actual plasmid fluctuations by time-averaging.

Reducing the problem of dampening increases the problem of expression noise—There are several potential solutions to the time-averaging problem, for example adding a degradation tag to the FP so that it adjusts more quickly to changing plasmid copy numbers, or placing the FP under an inducible promoter and inducing expression for a short time period. In both cases, short-lived FPs may not accumulate enough molecules to be reliably imaged, though this might be solved by using microscopy rather than FACS measurements. Slow fluorophore maturation also becomes a greater problem in both cases, but can in principle be solved by harvesting the cells, blocking further expression, and imaging cells after maturation has completed. However, both approaches also introduce other sources of heterogeneity: proteins that are actively degraded become susceptible to fluctuations in protease levels, while using inducible genes exposes proteins to ‘operator noise’ where some genes may be on when most are off and vice versa.

Another alternative is the use time-lapse microscopy to measure protein accumulation rates from time series, or to effectively erase the FP history by bleaching all existing fluorescent molecules. However, there is still a trade-off that cannot be resolved by any such method:

the goal is to prevent the protein from time-averaging out plasmid fluctuations, but that automatically prevents the protein from averaging out other sources of heterogeneity. Proteins that better track plasmids will thus also better track changes in for example mRNA levels, which increases the FP noise and masks real plasmid fluctuations. The noise contribution from the mRNA to the FP is indeed expected to be fairly small for stable proteins, but very large when considering the *rate* of making proteins, or the abundance of an unstable protein. Depending on the various sources of expression noise, there is thus a hard trade-off between reducing the problem of time-averaging and reducing the problem of expression noise, which could be an inherently unworkable trade-off.

Variances cannot be mathematically inferred without knowing the full dynamics of the system—A principally different solution to this problem is to attempt to infer the sources of heterogeneity by also tracking the expression of chromosomal genes in single cells (Wong Ng, Chatenay, Robert, & Piorer, 2010), ideally expressing a spectrally distinct FP in the same cells. The assumption is then that the noise in gene expression has similar dynamics for the chromosomal and plasmid genes, so that mathematical analysis can be used to back out the contribution from the plasmid. However, the analysis relies on unproven (and unlikely) assumptions about stochastic processes in cells. In fact, because plasmids and chromosomes have different replication control, the gene dosage should change with different dynamics, which in turn means that the time-averaging process is quantitatively different. That was in fact observed in the Løbner-Olesen experiment, where chromosomally expressed fluorescent proteins showed a multi-peaked distribution corresponding to the number of gene copies in individual cells, while even F plasmids with similar average gene dosage displayed single-peak distributions. We thus *know* that the processes are not equivalent in this way, and cannot assume that they are similar to infer the noise mathematically. Heavily relying on mathematically untested assumptions to infer the noise is then not all that different from the many studies that model the noise based on what is known about the control mechanisms. In a later paper partly from the same group (Ghozzi, Wong Ng, Chatenay, & Robert, 2010) they introduce the mathematical model used for estimating single cell plasmid copy numbers from data. They show that even if the assumptions were trusted they can only identify a broad interval of estimated variances, and thus only demonstrate consistency with the Løbner-Olesen paper, which is perhaps not so surprising since they use a similar method. Furthermore, the dual reporter method – where correlations between two reporters are used to track the origin of fluctuations – does not even work quantitatively in most linear systems where the reporters are identical and independent (Hilfinger & Paulsson, 2011) and much more extensive and rigorous multi-reporter systems would have to be designed to separate the various noise contributions.

Over the last five years, we have also tried many versions of these FP-expression approaches – moving to high-throughput microscopy, actively degrading or bleaching the fluorescent protein, measuring correlations between reporters etc. In each case controls showed that the distributions so obtained did not reflect the actual plasmid copy numbers. We still hope that some version of this method will work, by using better FPs as well as growth conditions in which we expect less noise in other components that affect expression rates.

2.2 Direct plasmid visualization

The conceptually most straightforward method of counting plasmids in cells is to fluorescently tag plasmid DNA and visually count spots. If this could be done in live cells, it would provide information not only about copy number distributions, but the full spatial and temporal dynamics: which plasmid copy replicates when and where. The challenge is that most plasmids naturally co-localize into diffraction-limited spots. A single spot in a conventional microscope image may then correspond to anything from a single plasmid

copy to tens of copies. Super-resolution microscopy techniques (Huang, Bates, & Zhuang, 2009) can in principle resolve objects as close as ten nanometers apart, but require bacterial cells to be fixed to repeatedly image a small fraction of photoactivatable fluorophores, and eventually pin point each light source by fitting distributions. As of yet it has not been used to reliably count molecules, and it is not known how closely plasmid copies co-localize within each diffraction limited spot. Another approach to deal with co-localization is using methods to quantify the total intensity of each spots to infer the number of copies. Here we discuss two visualization approaches and their associated problems: FISH and arrays of binding sites for fluorescently tagged DNA binding protein.

2.2.1 FISH—The standard method for visualizing polynucleic acids in cells is fluorescence in situ hybridization (FISH), where target sequences are specifically bound to labeled single-strand DNA probes. After fixing, hybridizing, and washing cells, labeled probes should ideally bind all the designated target molecules (no false negatives) and nowhere else (no false positives). However, given that the non-specific targets – not just DNA and RNA, but also other hydrogen bond donors – greatly outnumber the specific ones, the energetic differences between fully matched and mismatched interactions may be insufficient for accurate counting. There are many approaches to reduce the rate of false positives, such as increasing temperature, but they instead often increase the rate of false negatives. This trade-off between selective and quantitative binding is well studied in the context of micro-arrays where there are less competing interactions and more freedom to control relevant parameters. The same problem may be much more challenging to address *in situ* where the main strategy has been to carefully design primers.

One primer strategy that has been used for RNA is to design tens of small probes against each transcript (Raj, Peskin, Tranchina, Vargas, & Tyagi, 2006), thereby increasing the chance that multiple probes bind, while at the same time making sure that non-specific binding produces much lower fluorescence intensity. This method is useful for spatially separated targets, which is often the case in mammalian cells, since it greatly increases the intensity of a spot relative to the cytoplasmic background level.

However, when molecules are not spatially separated beyond the diffraction limit, as in the case of many plasmids, the number of spots does not represent the number of plasmids. Because quantitative super-resolution microscopy is not yet available for FISH, and we do not know how close the plasmid copies are, the intensity of each spot must then be quantified. The quantification can in turn be established in two ways: ideally the histogram of intensities per spot should display equidistant peaks corresponding to the integer number of molecules per spot, or a separate control experiment with single plasmid copies should show a relative standard deviation in intensity that is much smaller than the fold difference between the copy numbers that are being separated. This has so far not been possible, presumably due to heterogeneity in binding: when using a large number of probes it is practically impossible to evaluate the efficiency of each one, and the probe cocktail may thus contain several probes that bind poorly or non-specifically.

The other extreme FISH strategy is to use a single but highly optimized probe. No non-specific binding events are then allowed since they would be inseparable from the real events. However, since all spots would have the same expected intensity, it is also impossible to use the binding profile to test if the binding has gone to completion. Such strategies must therefore be accompanied by careful controls for example showing that the number of spots observed is insensitive to the probe concentration used.

An intermediate strategy, that was developed by Robert Singer's group in the first quantitative FISH studies for RNA counting (Zenklusen & Singer, 2010), is to use a handful

of probes against each target, and attach many fluorophores per probe. These probes are currently more expensive commercially and the total spot intensity is often lower than in the multi-probe FISH method above. However, having multiple fluorophores per probe makes it possible to separate between one and two probes using quantitative microscopy. The low number of probes further makes it feasible to evaluate each probe individually, both in the wildtype (here plasmid-containing) and the knock-out (here plasmid-free) strains, and thus discard any poorly working probes that would have increased signal-to-background for the total spot at the expense of decreasing the signal-to-noise in the intensity per spot. Furthermore, because it is possible to distinguish between spots with three versus two probes bound, the completion of binding can be more directly evaluated. Using multiple probes with different colors for the same plasmid can also improve the fidelity though accurate co-localization will probably require super-resolution methods to separate between plasmid copies.

FISH assays that do not provably meet the requirements for having narrow intensity distributions per plasmid copy must be dismissed as qualitative at best, and though they may give an indication of localization patterns, they cannot be used to count molecules. In fact, even if it was demonstrated that all plasmids that bind a probe of one color also bind a probe of another color, and even if the intensity distribution per spot had clear equidistant peaks corresponding to discrete plasmid copies, the problem remains that some plasmids may be entirely unavailable for binding – because of the crowded environment, proteins bound, the conformational state, or because the crosslinks introduced by fixation prevents binding altogether, for example creating microscopic pockets of the cell where probes cannot reach. The ideal control against such artifacts is to cross-validate the results with an independent counting method applied to the same single cells, rather than just the same populations. In lieu of that, it may be sufficient to compare the average copy number to bulk methods, keeping in mind that outliers contribute much more to standard deviations than they do to averages.

Plasmid FISH has been used by several groups, including Kurt Nordström (Weitao, Dasgupta, & Nordström, 2000), for measuring the localization of plasmids but has not been published as a means to count plasmids. Plasmid FISH suffers from all the issues mentioned above plus one more: the DNA strands must be separated to allow the probe to bind, usually by denaturation at high temperature. This could lead to interference since the second strand is fixed, as the whole cell is, and cannot diffuse away and may compete with the probes. We have also tried several FISH methods to count plasmids, and in our hands none of them have so far met these standards, though that may be a matter of protocol optimization.

2.2.3 Arrays of binding sites—Plasmids are most commonly visualized by inserting into the plasmid an array of DNA repeats that bind a specific protein tagged with a fluorescent reporter. Plasmids then generate fluorescent foci that are visible over the background, as long as the arrays are large enough. This method is attractive since it is done in live cells that provide much more information and where there are no risks of losing material or concerns about the partial penetration of probes.

The three most commonly used combinations of DNA binding sites and binding protein are lacO/lacI, tetO/tetR and ParB/parS. These arrays have been widely used to determine plasmid localization or to monitor the spatial dynamics of chromosomal genes. But when it comes to quantitatively *counting* the number of plasmids, at least the naïve method fails. Just as for FISH, plasmids often co-localize and the spots must either be resolved with super-resolution, or the intensities of the spots must be shown to quantitatively reflect plasmid copy numbers. Intensity quantification has unfortunately proven very difficult and

the intensity distributions have not displayed equidistant peaks corresponding to the integer number of copies, but rather tend to follow very broad and smooth distributions.

When using the ParB/parS system, quantification of the spots seems not to be an option at all, since after initial binding to parS, ParB polymerizes without a defined end point, though it may be possible to add a binding site for a strong DNA binding protein to block polymerization beyond a certain site. When working with lacO/lacI or tetO/tetR arrays, in order to ensure that the number of FPs bound per array does not fluctuate significantly due to individual binding and unbinding events, or varies with the cytoplasmic FP concentration, the FP concentration must be high enough so that the arrays are always saturated. In addition to the variable auto-fluorescence background of cells, this adds a high and variable background of cytoplasmic FPs that makes it harder to quantify the spot intensities. Furthermore, even if every binding site would bind an FP, only a fraction of the FPs will be fluorescent at any given time, which creates binomial fluctuations in the intensity per spot, or worse. Compounded with microscopy artifacts, this can randomize the spot intensity to the point where for example a spot with two plasmid copies is recorded as being brighter than a spot with three plasmid copies. The high local concentration of FPs at each array can also cause even reporters with very weak dimer formation tendencies to form large aggregates, much like avidity effects in antibody interactions. A protein may then dissociate from its binding site, and be replaced by a new protein from the cytoplasm, but the FP part of the dissociating fusion protein may still be bound to FPs at neighboring sites. As a result a plasmid can in fact have *more* FPs bound than it contains binding repeats, which would further contribute to the variation in fluorescence intensity per plasmid copy. In fact, if these aggregates dissociate from the plasmid but still stay bound together, they could mimic a plasmid focus, which can be misleading since no such aggregates would form in the plasmid-free cells. Our lab recently demonstrated FP induced clustering for a number of complex forming proteins in *Escherichia coli* (Landgraf, Okumus, Chien, Baker, & Paulsson, in press), as well as for LacI-FP fusions that bind to *lacO* repeats (unpublished).

The bound proteins also interfere with plasmid replication and movement (Sengupta, Nielsen, Youngren, & Austin, 2010). We have observed that complete binding to 240 or 120 repeat arrays prevents virtually all movement, replication or partitioning of the plasmid, so very careful controls would be required to show that partial binding does not also create significant effects. We have also evaluated the binding array approach with a variety of systems, microscopy methods, FPs and plasmids, and applied them to cells where we know there is only a single copy of the plasmid, to directly estimate the heterogeneity in the read-out. So far, and in our hands, none of the methods have met the necessary conditions for reliable quantitative counting. If such a method were to work in the future we expect it to involve highly monomeric reporter proteins (Landgraf, Okumus, Chien, Baker, & Paulsson, in press), not so many binding sites that the complex interferes with plasmid dynamics but enough to reliably detect the spot above background, possibly super-resolution methods to resolve individual plasmid copies, and controls that demonstrate that the distribution of intensity per focus has equidistant peaks corresponding to the underlying number of plasmid copies.

2.3 PCR based methods

The classic approach to detect specific DNA targets is PCR amplification. But standard PCR methods are not quantitative: the amplification process is exponential and after a few amplification cycles, the amount of product correlates poorly with the amount of initial material. The accuracy of PCR has been improved in many ways, introducing new polymerases, following the process in real time, using better dyes etc. Here we discuss the opportunities and challenges associated with quantitative and digital PCR.

2.3.1 qPCR—Quantitative PCR (qPCR) is now widely used to measure RNA and DNA levels but for many applications the term is a misnomer. At low molecule numbers probabilistic chemistry can introduce errors that then are amplified in subsequent cycles: with a single initial copy, a failure to replicate in the first cycle would produce a two-fold error that, due to the exponential nature, would persist through the cycles. As all methods that rely on DNA hybridization, there is also a trade-off between hybridizing to all target molecules (no false positives) and only hybridizing target molecules (no false negatives). Errors are then hard to avoid and it is therefore standard to run at least triplicates per sample, which is impossible when working with single cells.

qPCR is also sensitive to several external factors and is therefore mostly used for reporting relative amounts between different samples within one experiment, where converting to absolute numbers requires a calibration curve with known amounts of DNA or RNA in the relevant range. The calibration itself must then be established separately, usually with UV measurements for serial dilutions which have their own sources of errors.

Single cell analyses introduce several additional problems. For typical plasmid copy numbers, ranging from one copy to a few tens of copies, it is difficult to establish reliable calibration curves. Even if such curves could be confirmed, calibration curves often account for the fraction of false negatives. For example, if only 50% of molecules are detected by PCR, that can simply be included into the calibration. At low numbers however, a 50% chance per molecule of not being detected will cause binomial fluctuations around the average, sometimes detecting all copies and sometimes none. Such effects cannot be corrected for via calibration and can introduce large errors. Adding pre-amplification steps before running qPCR introduces the same type of amplification errors as qPCR itself.

In single cell assays it can also be challenging to conserve the full sample volume, which may require that individual cells are lysed inside the reaction tube, which in turn requires methods for accurately inserting exactly one cell per tube, and ideally methods to measure the size of each cell. Generating distributions in turn requires hundreds of samples, and ideally many more to look for equidistant peaks in the histogram. To our knowledge the few single cell studies based on qPCR have therefore focused on high copy RNAs, and our own attempts to quantify plasmid expressed RNAs with qPCR were far from quantitative. We do not know if this is due to non-quantitative PCR or to the fact that the RNA still only partially reflects the plasmid copy number in individual cells, see section 2.1

2.3.2 Digital PCR—A strategy better suited for the low copy regime of plasmids is digital PCR. The idea is to dilute the sample into many wells such that most wells contain zero copies of the DNA or RNA of interest, a few wells contain one copy, and almost no wells contain more than one. Subsequent PCR should then yield some products in the wells with one copy and not otherwise, and hence the number of wells with signal should correspond to the number of DNA or RNA molecules in the original sample.

This concept was introduced in 1999 (Vogelstein & Kinzler, 1999), and variants have recently become commercially available in different platforms – including the nano-wells system by Fluidigm and more recently the droplet-based platform from Quantalife and RainDance. Digital PCR side-steps the main problem of qPCR of errors in the first cycles: if it is sufficient to separate between having template and not having template the PCR reactions do not need to be as quantitative.

However, though conceptually straightforward it is practically difficult to implement without introducing artifacts. In addition to detection problems when a single specific target competes with a huge number of non-specific targets, which could be partially addressed by

amplifying many targets per plasmid, the main challenge is perhaps the cell handling. First precisely one cell must be placed in a lysing well, ideally after recording its size. In our experiences FACS sorting is unreliable for both – even at low sample rates, it sometimes picks two cells – and light scatter is a very poor measure of size of bacteria. The cell must then lyse completely and mix with the PCR cocktail, and the multiplasmid foci must disintegrate into individual copies, which must be validated independently. Then the complete volume of the single cell lysate must be subdivided into such a large number of subvolumes that two plasmids are unlikely to end up in the same subvolume. The current commercial devices do not include on-chip cell handling, and also have a dead volume that never enters the PCR reactions. They therefore do not claim to have the resolution needed for plasmid counting (our local core facility set up with the Fluidigm system advised that the assay will not be quantitative unless we have hundreds of molecules per cell).

In addition to evaluating commercial digital PCR systems, we are completing a custom microfluidic device to deal with all these challenges, to count plasmid copy numbers in single cells without need for cloning.

3. Replication arrest methods

The perhaps most promising assay to determine plasmid copy numbers accurately is to block plasmid replication abruptly and completely without interfering with chromosomal replication or cell growth. As each cell grows into a micro-colony, plasmids eventually segregate into separate cells. This reduces the problem of counting individual plasmid copies in cells to counting plasmid-containing cells in the population descending from the founder cell when replication was arrested. The final detection can be done in many ways: expressing a resistance gene and selecting for plasmids, expressing GFP and separating between cells that have some versus no GFP rather than quantifying fluorescence levels, and tagging plasmids using FP-repressor fusions and arrays of binding sites, now without the complication that each spot can contain multiple copies. The approach thus combines several assays above, with the advantage of a digital readout, exploiting the natural cell growth and plasmid segregation to minimize material loss, and expressing plasmid encoded genes to achieve higher detection accuracy than FISH or PCR.

The main challenge for these methods is the replication arrest, which must be immediate and absolute. There are many replication arrest methods reported in the literature and we have carefully and systematically evaluated pSC101 and F plasmids with temperature sensitive replication proteins, and R1 with the sudden induction of the replication inhibitor CopA. The seemingly straightforward methods described in the literature are not quantitative however. For example, the complete replication arrest observed for some pSC101 strains is not immediate, but allows for partial replication before coming to a complete halt, which distorts the distributions. This is invisible in the classic control that the number of plasmid containing cells post replication arrest approaches a perfectly stable plateau determined by the average plasmid copy number. Increasing the temperature to the point where plasmids stop replicating immediately will instead cause cell death that interferes with the counting assay. Replication arrest of plasmid R1 suffered similar problems but for different reasons.

We have worked extensively to optimize the arrest methods to count plasmid pSC101 and R1 with high accuracy in single cells, as well as their partitioning mechanisms at cell division, and believe they can be made to work accurately. However, the approach is very labor intensive to perform with appropriate controls, and still only provides a snapshot of the copy number at the time of the replication arrest. It is also difficult to extend to additional plasmid strains without extensive controls and likely modifications – it may even be difficult to extend to different growth conditions without repeating all controls. For these reasons we

are in parallel pursuing many of the approaches described in sections 2 and 3, as well as several others.

4. Conclusions

We analyzed a range of potential methods for measuring plasmid copy numbers in single cells, and the challenges that must be overcome to make them quantitative. All methods produce an estimated distribution, but without extensive optimization and evaluation, we have strong reason to believe that the reported fluctuations have little to do with the actual plasmid copy numbers. The methods that we believe are the most promising, based on replication arrest, are very labor intensive, require extensive controls and optimized procedures, and are difficult to extend to more types of plasmids. We therefore continue to work to develop alternative methods that provide more information about the underlying process, allow greater sampling, and that can be extended to a broader range of plasmids. We believe these methods will be fully developed in the next few years, by us or others, and that the dynamics of plasmid control mechanisms then can be studied quantitatively in their natural context: the individual cell.

Acknowledgments

We thank Kurt Nordström and Per Malkus for discussions. This work was supported by the US National Institutes of Health (Grant GM081563-01A). S.T. thanks the Human Frontier Science Program for a cross-disciplinary fellowship.

References

Becskei A, Serrano L. Engineering stability in gene networks by autoregulation. Nature. 2000; 405(6786):590–593. [PubMed: 10850721]

Ghozzi S, Wong Ng J, Chatenay D, Robert J. Inference of plasmid-copy-number mean and noise from single-cell gene expression data. Phys Rev E. 2010; 82:051916.

Hilfinger A, Paulsson J. Separating intrinsic from extrinsic fluctuations in dynamic biological systems. Proc Natl Acad Sci USA. 2011; 108 (29):12167–12172. [PubMed: 21730172]

Huang B, Bates M, Zhuang X. Super resolution fluorescence microscopy. Annu Rev Biochem. 2009; 78:993–1016. [PubMed: 19489737]

Landgraf D, Okumus B, Chien P, Baker TA, Paulsson J. Segregation of molecules at bacterial cell division reveals native protein localization. Nat Biotechnol. in press.

Lestas I, Vinnicombe G, Paulsson J. Fundamental limits on the suppression of molecular fluctuations. Nature. 2010; 467(7312):174–178. [PubMed: 20829788]

Løbner-Olesen A. Distribution of minichromosomes in individual Escherichia coli cells: implications for replication control. EMBO J. 1999; 18(6):1712. [PubMed: 10075940]

Nordström K, Molin S, Light J. Control of replication of bacterial plasmids: genetics, molecular biology, and physiology of the plasmid R1 system. Plasmid. 1984; 12:71–90. [PubMed: 6209739]

Raj A, Peskin CS, Tranchina D, Vargas DY, Tyagi S. Stochastic mRNA synthesis in mammalian cells. PLoS biol. 2006; 4(10):e309. [PubMed: 17048983]

Sengupta M, Nielsen HJ, Youngren B, Austin S. P1 plasmid segregation: accurate redistribution by dynamic plasmid pairing and separation. J bacteriol. 2010; 192(5):1175–1183. [PubMed: 19897644]

Vogelstein B, Kinzler KW. Digital PCR. Proc Natl Acad Sci USA. 1999; 96:9236–9241. [PubMed: 10430926]

Weitao T, Dasgupta S, Nordström K. Plasmid R1 Is Present as Clusters in the Cells of Escherichia coli. Plasmid. 2000; 43:200–204. [PubMed: 10783298]

Wong Ng J, Chatenay D, Robert J, Piorer MG. Plasmid copy number noise in monoclonal populations of bacteria. Phys Rev E. 2010; 81:011909.

Zenklusen D, Singer RH. Analyzing mRNA expression using single mRNA resolution fluorescent in situ hybridization. Method enzymol. 2010; 470:641–659.

Research highlights

- Plasmid copy numbers have not been accurately measured in single cells
- Most current methods confound plasmid fluctuations with other stochastic processes.
- In some cases, measurement errors mask plasmid fluctuations.
- Some problematic methods could be rectified by modifications to existing protocols