

Simulating single-cell metabolism using a stochastic flux-balance analysis algorithm

David S. Tourigny,^{1,2,*} Arthur P. Goldberg,³ and Jonathan R. Karr³

¹Irving Institute for Cancer Dynamics, Columbia University, New York, New York; ²School of Mathematics, University of Birmingham, Birmingham, United Kingdom; and ³Icahn Institute for Data Science and Genomic Technology, and Department of Genetics and Genomic Sciences, Icahn School of Medicine at Mount Sinai, New York, New York

ABSTRACT Stochasticity from gene expression in single cells is known to drive metabolic heterogeneity at the level of cellular populations, which is understood to have important consequences for issues such as microbial drug tolerance and treatment of human diseases like cancer. Despite considerable advancements in profiling the genomes, transcriptomes, and proteomes of single cells, it remains difficult to experimentally characterize their metabolism at the genome scale. Computational methods could bridge this gap toward a systems understanding of single-cell biology. To address this challenge, we developed stochastic simulation algorithm with flux-balance analysis embedded (SSA-FBA), a computational framework for simulating the stochastic dynamics of the metabolism of individual cells using genome-scale metabolic models with experimental estimates of gene expression and enzymatic reaction rate parameters. SSA-FBA extends the constraint-based modeling formalism of metabolic network modeling to the single-cell regime, enabling simulation when experimentation is intractable. We also developed an efficient implementation of SSA-FBA that leverages the topology of embedded flux-balance analysis models to significantly reduce the computational cost of simulation. As a preliminary case study, we built a reduced single-cell model of *Mycoplasma pneumoniae* and used SSA-FBA to illustrate the role of stochasticity on the dynamics of metabolism at the single-cell level.

SIGNIFICANCE Because of fundamental challenges limiting the experimental characterization of metabolism within individual cells, computational methods are needed to help infer the metabolic behavior of single cells from information about their transcriptomes and proteomes. In this article, we present a stochastic simulation algorithm with flux-balance analysis embedded, the first systematic framework, to our knowledge, for modeling the stochastic dynamics of single cells at the level of genome-scale metabolic reaction networks. We provide a robust and efficient algorithm for simulating stochastic simulation algorithm with flux-balance analysis embedded models and apply it to a case study involving the metabolism, RNA, and protein synthesis and turnover of a single *Mycoplasma pneumoniae* cell.

INTRODUCTION

Describing the phenotypic behavior of single cells is critical for a deeper understanding of biological tissues, organisms, and populations. Recent experimental advances are driving a data explosion in systems biology by enabling researchers to profile multiple dimensions of single cells, including their genome, transcriptome, and proteome (1–3). Such single-cell measurements can yield information about thousands of individual cells in a single experiment. This can provide insights on intracellular function and the role of intercellular heterogeneity in a variety of biological systems ranging

from microbial populations (4,5) to human diseases such as cancer (6,7).

Although metabolism is a key aspect of cellular physiology, methodologies for probing the metabolism of single cells are comparatively immature (8–12). This presents a barrier to studying a variety of phenomena such as metabolic reprogramming in tumors, which is now understood to be a central hallmark of cancer (13,14). It is challenging to experimentally probe the metabolism of single cells due to low abundances of many metabolites, the compartmentalization of eukaryotic cells, and the wide structural diversity of metabolites (10). Moreover, faster timescales of enzymatic reactions make measuring the dynamics of metabolism more difficult than measuring that of DNA replication or gene expression. These experimental obstacles to studying single-cell metabolism necessitate the

Submitted June 13, 2021, and accepted for publication October 26, 2021.

*Correspondence: dst2156@columbia.edu

Editor: Ramon Grima.

<https://doi.org/10.1016/j.bpj.2021.10.038>

© 2021 Biophysical Society.



development of computational techniques that can infer the metabolism of single cells from other information, such as single-cell transcriptomics or proteomics and population-level metabolic data (15).

Although our current capabilities to probe and model the metabolism of single cells are limited, considerable attention has been devoted to the metabolism of cellular populations. For example, metabolic network modeling has achieved a great deal of success combining limited experimental data and computational simulation to study gene essentiality and guide metabolic engineering (16–18). Extensions of these population-based metabolic network modeling frameworks, such as dynamic metabolism expression models (19) and dynamic enzyme-cost flux-balance analysis (FBA) (20), have also been developed to incorporate the dynamics of gene expression. However, these deterministic approaches cannot capture effects that are relevant to the metabolism of single cells, in which stochasticity is understood to play a pivotal role (21,22). To date, there have only been a handful of attempts (see (23) and references therein) to extend metabolic network modeling to the single-cell regime. These efforts have mainly focused on integrating single-cell transcriptomics data with FBA in the context of cancer. Critically, these studies have not addressed the temporal dynamics of single cells, which is a key feature of their behavior.

The existing approaches to stochastic modeling of single-cell metabolism can be organized into two categories: 1) (semi-)analytical treatment of rigorous models of individual metabolic pathways involving the expression of one (24,25) or several (26,27) enzymes catalyzing a handful of metabolic reactions, or 2) empirical simulation of whole-cell models (28,29) involving hybrid methods that combine ordinary differential equations (ODEs), particle-based stochastic simulation algorithms (SSAs (30,31)), and dynamic FBA (DFBA (32)). Although the shared aim of these approaches is to relate single-cell behavior to genotype, the two categories fall at opposite ends of a wide spectrum: whole-cell modeling aims to accommodate as much detail as possible, but no framework is yet available to rigorously simulate entire cells. On the other hand, analytical methods that can provide a mechanistic understanding of small pathways are intractable for entire cells. Furthermore, existing exact, approximate, and hybrid stochastic simulation methods for large-scale biochemical networks (reviewed in (33)) remain a long way from applicability to single-cell metabolism because, unlike FBA, they still rely on explicitly encoding the dependence of reaction rate functions on enzymatic kinetic parameters. Here, we introduce a stochastic extension of FBA that we call SSA with FBA embedded (SSA-FBA), which enables systematic systems-scale simulations of the metabolism of single cells. We believe SSA-FBA is a powerful computational tool for simulating stochastic dynamics of metabolic network models at the single-cell level.

The remainder of this article is organized as follows. We first outline the concepts of SSA-FBA and relate it to a formal description of single-cell metabolism based on the chemical master equation (CME). Next, we compare exact and approximate implementations of SSA-FBA and then introduce an advanced algorithm that significantly improves the efficiency of exact simulations. As a case study, we subsequently present simulation results of an SSA-FBA model of a reduced single *Mycoplasma pneumoniae* cell representing 505 reactions in which, as an illustrative example, we can explore the consequences of stochasticity for the dynamics of ATP production and consumption at the single-cell level. We conclude the article with a discussion of the main results and directions for future work.

MATERIALS AND METHODS

We implemented an SSA-FBA simulation package in C++ and Python. Further details on SSA-FBA, its implementation, and the case study can be found in the [Supporting materials and methods](#), and all code and data are freely available open-source at <https://gitlab.com/davidtourigny/single-cell-fba>.

RESULTS

SSA-FBA

Stochasticity in the metabolism and growth of single cells is generally believed to emerge primarily from fluctuations in enzyme expression levels (21,22,24,25,34) (see Fig. 1, which provides an overview of sources and consequences of stochasticity at the single-cell level). Because of the relatively high copy numbers of most metabolites, metabolism is thought to have little intrinsic stochastic variation compared with gene expression (21). This observation motivates SSA-FBA as an appropriate framework for modeling single-cell metabolism, in which the dynamics of reaction fluxes internal to a metabolic network are captured deterministically by FBA, whereas SSA is used to model changes in the copy numbers of enzyme molecules and metabolites that are produced or consumed on the periphery of the metabolic network.

The combined single-cell network of M metabolic and enzyme expression reactions can be captured by the CME (35)

$$\frac{dP(\mathbf{n})}{dt} = \sum_{j=1}^M a_j(\mathbf{n} - \mathbf{S}_j)P(\mathbf{n} - \mathbf{S}_j) - a_j(\mathbf{n})P(\mathbf{n}), \quad (1)$$

where \mathbf{n} is an N -dimensional vector for counts of N chemical species, $a_j(\mathbf{n})$ is the propensity value of reaction j given \mathbf{n} , and \mathbf{S}_j is the stoichiometry of reaction j . The probability density function $P(\mathbf{n})$ describes the probability of the system to occupy state $\mathbf{n}(t)$ at time t . Because CMEs are usually too complex to be solved directly, investigators often use SSA to

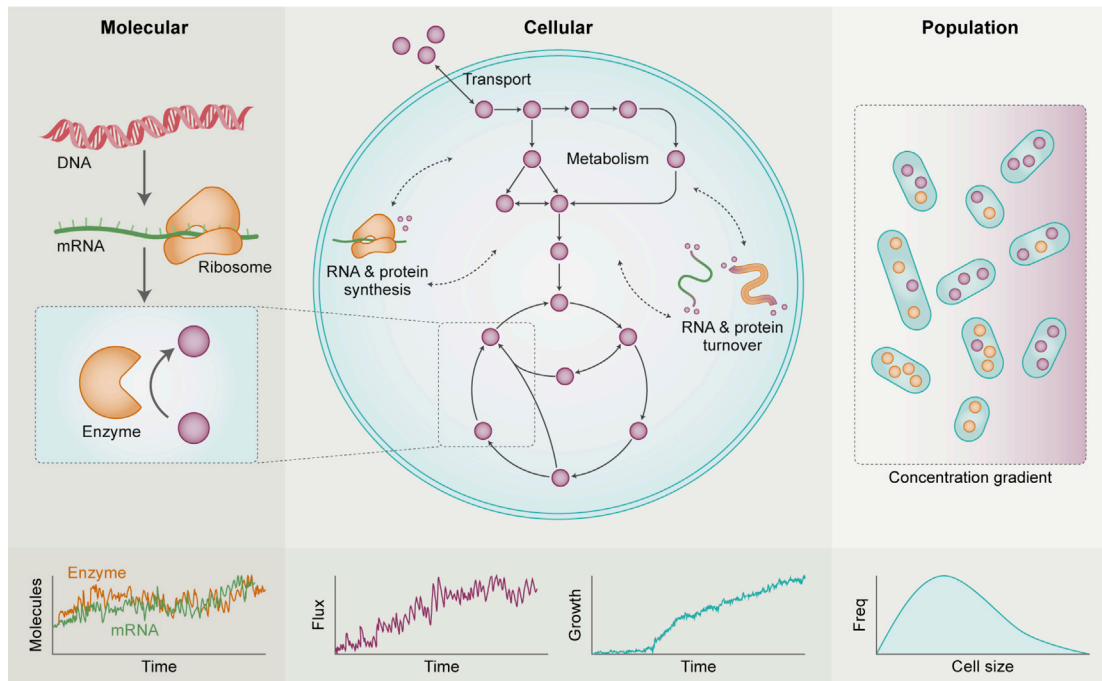


FIGURE 1 The metabolic behavior of single cells is influenced by the spatiotemporal dynamics of metabolites and enzymes at multiple scales. Temporal variation in the metabolic behavior of single cells is primarily driven by fluctuations in gene expression and amplified or attenuated by their metabolic networks. For example, the structure of the reaction network combined with fluctuating enzyme and metabolite levels determine intracellular metabolic fluxes that in turn govern the dynamics of energy supply and molecular compositions of individual cells. The latter determines a total cell mass, which, assuming typical cell density, can be used to infer cell volume or size. The resulting stochastic nature of these properties at the single-cell level drives phenotypic heterogeneity modeled as a distribution over random variables at the population level, which can also be influenced by external factors such as environmental conditions. To see this figure in color, go online.

sample trajectories through the state space of the CME (30,31). One limitation of SSA is its computational cost, which is particularly acute for large chemical reaction networks. Various performance enhancements have therefore been developed to improve the run time of SSA. Because the average counts of enzymes are many times lower than those of metabolites and metabolic reactions operate on a much faster timescale than that of gene synthesis and degradation, several researchers have used stochastic quasi-steady-state assumptions (stochastic generalizations of the quasi-steady-state assumption in DFBA (32)) to approximate SSA trajectories at a lower computational cost. These methods correspond to a reduction of the CME (1) on the basis of timescale separation or molecular abundances (see [Supporting materials and methods](#), Section S1; (36–39)).

Although the above timescale or abundance separation methods could in principle be used to simulate single-cell metabolism, investigators rarely have sufficient data to construct detailed dynamical models of metabolism in single cells because of the experimental challenges outlined in [Introduction](#). Instead, we realized that the constraint-based formalism of (D)FBA (16,17,32) could be used to identify a numerical solution to the deterministic quasi-steady-state conditions of the metabolic reaction network

with a small amount of data at a modest computational cost. In short, SSA-FBA embeds an FBA model into SSA, analogous to the way that an FBA model is embedded into a system of ODEs in DFBA. Marginalization of the CME over copy numbers of metabolites internal to the metabolic reaction network motivates describing the dynamics of macromolecules and metabolites external to the metabolic reaction network using SSA but when propensity values for metabolic reactions are obtained by solving the embedded FBA problem. In turn, the embedded FBA problem and therefore the resulting propensity values depend on the counts of macromolecules and external metabolites, which implies the embedded FBA problem must be dynamically updated and resolved for metabolic propensity values over the course of simulation. This embedding enables SSA-FBA to simulate the stochastic behavior of the metabolism of single cells analogous to how DFBA enables deterministic simulation of the dynamic behavior of the metabolism of populations.

First, SSA-FBA separates reactions into three mutually disjoint subsets based on whether the species that participate in the reaction are defined to be internal or external to the metabolic reaction network ([Fig. 2 A](#)). The three subsets of reactions in an SSA-FBA model are as follows:

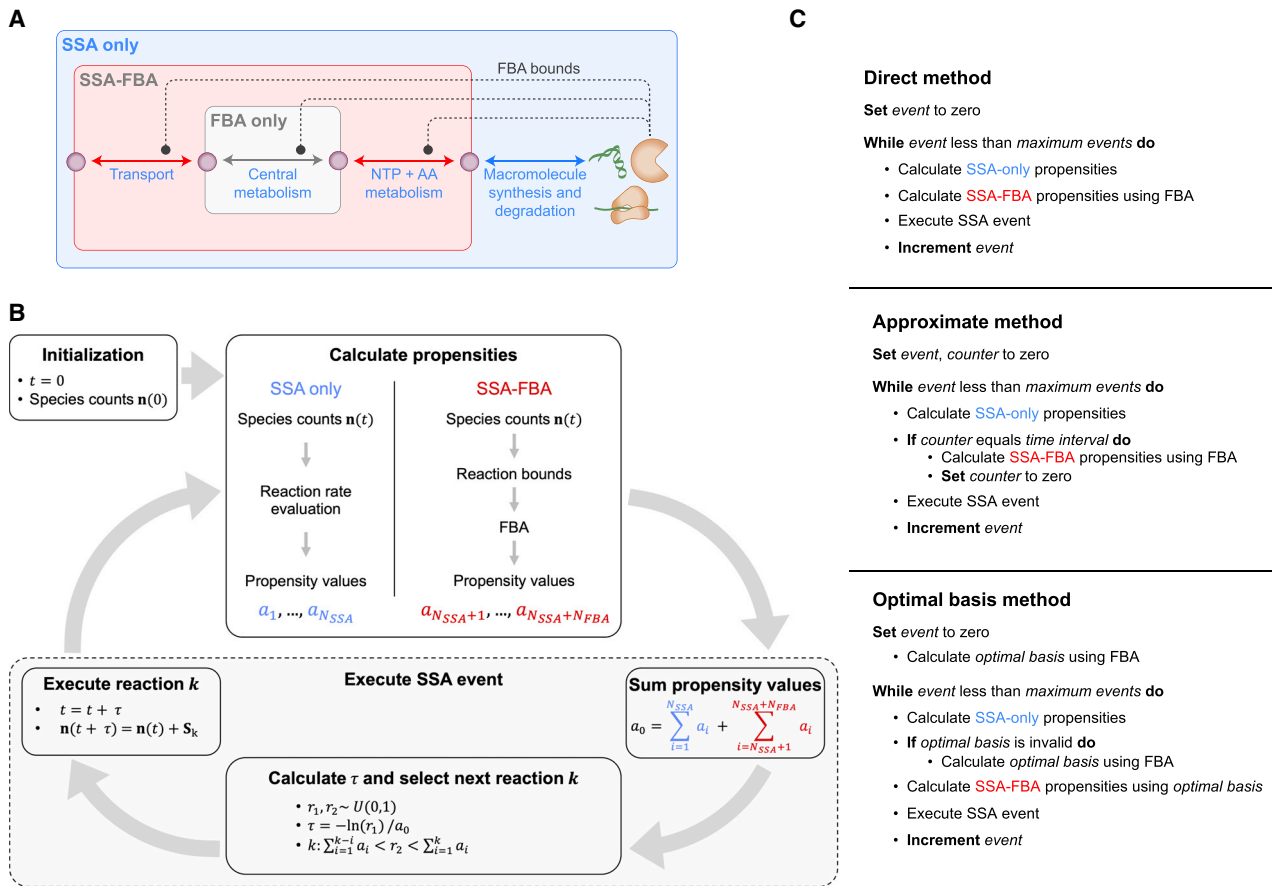


FIGURE 2 Overview of SSA-FBA model structure and simulation algorithms. (A) SSA-FBA separates reactions into SSA-only (blue), SSA-FBA (red), and FBA-only (gray) subsets based on the species that participate in each reaction. In this example that models the coupling of transmembrane transport and central metabolism to macromolecule synthesis and degradation via amino acid (AA) and NTP metabolism, metabolites in the central carbon metabolism are internal to the metabolic reaction network. Gene products such as enzymes and transporters contribute to the rates or propensity values of FBA and SSA-FBA reactions, and this is modeled through FBA bounds on these reactions (dotted lines). (B) After initialization, each iteration of an SSA-FBA simulation first calculates the propensity values for SSA-only reactions (blue) by direct evaluation of reaction rate functions (as in Gillespie’s original algorithm (30,31)) and calculates optimal propensity values for SSA-FBA reactions (red) using an optimal solution of the FBA model in which current species counts are used in reaction bounds. Next, both types of propensity values are combined to select the next reaction and its execution time. Finally, species counts are updated according to the stoichiometry of the selected reaction. (C) Pseudocode for the direct, approximate, and optimal basis SSA-FBA simulation methods. The approximate method speeds up SSA-FBA simulation by only calculating SSA-FBA propensity values at a fixed time interval (represented by Δ_{event} in the main text). The optimal basis method speeds up SSA-FBA without approximation by only calculating SSA-FBA propensity values when the optimal basis has changed. To see this figure in color, go online.

- 1) FBA-only reactions: reactions that interconvert among the internal species.
- 2) SSA-only reactions: reactions that interconvert among the external species.
- 3) SSA-FBA reactions: reactions that convert between the internal and external species.

For example, the SSA-only reactions may correspond to the synthesis and degradation of gene products, whereas the FBA-only and SSA-FBA reactions may include reactions involved in transmembrane transport and central metabolism.

Because external species, such as enzymes and transporters, often influence the rates or propensity values of metabolic reactions, the FBA-only and SSA-FBA reactions

can be constrained or bounded by the dynamic counts of external species (Fig. 2 A). Mathematically, this can be expressed as the following linear programming (LP) problem:

$$\begin{aligned} &\text{maximize: } z = \mathbf{c} \cdot \mathbf{a}_{FBA} \\ &\text{subject to: } \mathbf{S}_{FBA} \cdot \mathbf{a}_{FBA} = 0, \mathbf{l}(\mathbf{n}_{ex}) \leq \mathbf{a}_{FBA} \leq \mathbf{u}(\mathbf{n}_{ex}), \end{aligned} \quad (2)$$

where \mathbf{a}_{FBA} is a vector containing the propensity values of FBA-only and SSA-FBA reactions, \mathbf{S}_{FBA} is the submatrix of \mathbf{S} encoding the stoichiometry of the metabolic reaction network, and $\mathbf{l}(\mathbf{n}_{ex})$, $\mathbf{u}(\mathbf{n}_{ex})$ are bounds that depend on the counts of external species, \mathbf{n}_{ex} (see Supporting materials and methods, Section S1 for extended discussion). The

functional forms of bounds $\mathbf{l}(\mathbf{n}_{ex})$, $\mathbf{u}(\mathbf{n}_{ex})$ rely on the way that different reactions are represented in the model: for example, bounds for transport reactions could depend on species counts analogously to the way that exchange flux bounds can depend on extracellular substrate concentrations in DFBA (32), whereas propensity values of enzymatic reactions can be bounded in proportion to the abundances of intracellular enzyme molecules (so-called enzyme capacity constraints, e.g., (20,40)). These considerations were implemented in the whole-cell model described in (28), and also, the reduced model of *M. pneumoniae* was presented as an SSA-FBA case study in this article (see [Supporting materials and methods](#), Section S3). The coefficient vector \mathbf{c} is chosen to reflect a biologically relevant objective, such as maximizing the rate of production of the metabolites required for growth. Solving the LP problem (2) for a given instance of the external species counts vector \mathbf{n}_{ex} returns an optimal set of propensity values for FBA-only and SSA-FBA reactions, whereas propensity values for the SSA-only reactions are calculated from \mathbf{n}_{ex} by direct evaluation of a reaction rate function as in SSA (30,31).

Next, as outlined in [Fig. 2 B](#), SSA-FBA combines propensity values for N_{FBA} SSA-FBA reactions obtained from an optimal solution to the LP problem (2) with the propensity values of N_{SSA} SSA-only reactions to determine the next reaction event according to Gillespie's original algorithm (30,31) (although this step can also be replaced by more advanced methods such as the Next Reaction Method (41) or possibly incorporated into τ -leaping-based approximation algorithms (42)). Because FBA-only reactions do not affect the counts of external species, their propensity values are not required at this stage. Finally, execution of either an SSA-only or SSA-FBA reaction updates the species counts vector, which is in turn used to update the bounds of both FBA-only and SSA-FBA propensity values in (2). It is important to highlight that SSA-FBA is not restricted to cases in which the metabolic portion of a model is represented by an LP problem, as framed here, but can be extended more generally to scenarios in which nonlinear optimization problems used to calculate SSA-FBA propensity values are embedded within SSA. For example, inclusion of thermodynamic constraints can further restrict the possible set of SSA-FBA propensity values and results in a nonlinear or mixed-integer LP problem (e.g., (43)); however, the efficient simulation algorithm presented in the next section cannot necessarily simulate such SSA-FBA models because it depends on the optimal basis structure of LP problems.

In summary, SSA-FBA embeds FBA within SSA by using an LP problem to predict an optimal propensity value for each metabolic reaction using the counts of species predicted by SSA. This bidirectional coupling of SSA and FBA is analogous to the coupling of FBA and ODE integration in DFBA (32). One challenge of both DFBA and SSA-FBA is that solutions to (2)—and therefore the

optimal propensity values of SSA-FBA reactions—are not necessarily unique (44). However, it turns out to always be possible to formulate a lexicographic version of the LP problem by including multiple biological objective functions to be optimized sequentially to guarantee uniqueness and retain compatibility with the implementation described in the next section (see also (45)), although construction of a lexicographic LP is, in general, not unique either and imposes additional biological assumptions. It is important to highlight that unlike variants of SSA that approximate trajectories of the CME (1) to reduce the computational overhead (36–39), the goal of SSA-FBA is only to approximately predict the metabolic dynamics of single cells in the absence of detailed kinetic information about each species and reaction (16,17). The modeler can control the degree of this approximation by how they choose to partition reactions into the three subsets (FBA only, SSA only, or SSA-FBA). As a rule of thumb, we recommend that modelers distinguish between internal and external species (and hence partitioning of reactions) based loosely on the timescales of the species, but often, the distinction between SSA-FBA and FBA-only reactions will be largely determined by the structure of the metabolic model and the metabolites that are considered substrates for the gene expression (SSA-only) reactions. Further guidelines for encoding SSA-FBA models and additional information about the relationship between SSA-FBA and the CME are outlined in [Supporting materials and methods](#), Section S1.

Efficient implementation of SSA-FBA

As described above, SSA-FBA iteratively 1) uses the counts of the species and reaction rate laws to calculate the propensity values for each SSA-only reaction, 2) uses the counts of the species in FBA to calculate the propensity values of each SSA-FBA reaction, 3) uses SSA to select the next reaction, and 4) updates the species counts based on stoichiometry of the selected SSA-only or SSA-FBA reaction. The direct method for performing an SSA-FBA simulation in this way is summarized in the top row of [Fig. 2 C](#). At first glance, the direct method appears to be computationally expensive for larger models because it appears to require solving one FBA problem per SSA execution event.

One way to reduce the computational cost associated with the direct SSA-FBA simulation method is to approximate the SSA-FBA propensity values by updating them less frequently, assuming a time interval (Δ_{event}) across which an optimal solution to the embedded FBA problem does not change appreciably. When the end of the interval is reached, the constraints of the embedded FBA problem can be updated using the current counts of species and resolved for a new set of SSA-FBA propensity values used to parameterize integration over the next interval. This approximate SSA-FBA simulation method (*center row* in [Fig. 2 C](#)) is analogous to the simplest implementation

of DFBA (32) that has been used in the simulation of whole-cell models (28,29). To evaluate the performance of the direct and approximate simulation methods, we designed a toy model that contained only two SSA-only reactions (R0 and R1), one SSA-FBA reaction (R2, growth rate of *M. genitalium*) and one variable FBA bound (on the oxygen transport reaction). The toy model contained three species, S_0 , S_1 , and S_2 , and is represented by the reaction schema



In this schema, the propensity value of the SSA-FBA reaction R2 is the current optimal value of the growth rate objective function calculated by FBA using the metabolic network model of *M. genitalium* from (28). The R2 propensity value thus depends indirectly on species S_1 , which bounds the maximal rate of the oxygen uptake reaction in the *M. genitalium* metabolic network model based on the functional form $u(S_1) = S_1/10$. The optimal value of the objective function of the *M. genitalium* metabolic network model obtained by solving the associated FBA problem at this substrate level (corresponding to a given maximal rate of oxygen uptake) is then used as propensity value for reaction R2.

We began simulations with initial species counts $S_0 = 1000$, $S_1 = 0$, and $S_2 = 0$ and simulated the toy model across 3000 reaction execution events using either the direct or approximate SSA-FBA simulation method. In particular, we used the toy model to explore how the choice of Δ_{event} affects results of SSA-FBA simulations using the approximate method. When Δ_{event} is small, the variances of approximate SSA-FBA simulations are comparable with those obtained using the direct method (Fig. 3, A and B), and increasing Δ_{event} increases this variance (Fig. 3 C). The toy model highlighted a shortcoming of the approximate SSA-FBA simulation method that trajectories occasionally (particularly for larger Δ_{event} , as visible in Fig. 3 C, but

observed for all Δ_{event} tested) display numerical instability and diverge from realistic values toward the latter end of a simulation. This type of numerical instability also arises during direct integration of many DFBA models (45) but is particularly problematic for SSA-FBA because its stochastic nature means the pathology cannot be conclusively ruled out on the basis of trial simulations (for example, to establish optimal size of Δ_{event}). The origin of the problem is that the embedded FBA problem can become infeasible within an integration interval, which induces a closed domain of definition for the dynamic system (45).

We therefore searched for a strategy to reduce the computational cost of SSA-FBA simulation without the loss of accuracy and numerical stability associated with approximation. Motivated by a recent method for numerical integration of ODEs with embedded LP problems (45), we developed a more efficient algorithm that leverages the fact that FBA problems only need to be resolved when their optimal basis changes (outlined in bottom row of Fig. 2 C; full details in Supporting materials and methods, Section S2), which typically occurs much less frequently than after the execution of every SSA reaction because of the different timescales of metabolic and gene expression reactions. In this approach, only the validity of the current optimal basis being used to calculate optimal propensity values for SSA-FBA reactions must be established after each reaction execution event, which in the context of most SSA-FBA models, is computationally cheaper than solving the embedded LP problem each time molecular species counts are updated. The result of the optimal basis simulation method remains mathematically equivalent to that of the direct SSA-FBA simulation method but is obtained substantially faster because it requires solving many fewer LP problems. Fig. 3 D displays representative results of the comparisons (described in Supporting materials and methods, Section S2) between the optimal basis and direct (equivalent to the approximate with $\Delta_{event} = 0$) SSA-FBA simulation methods, showing that the former improves run

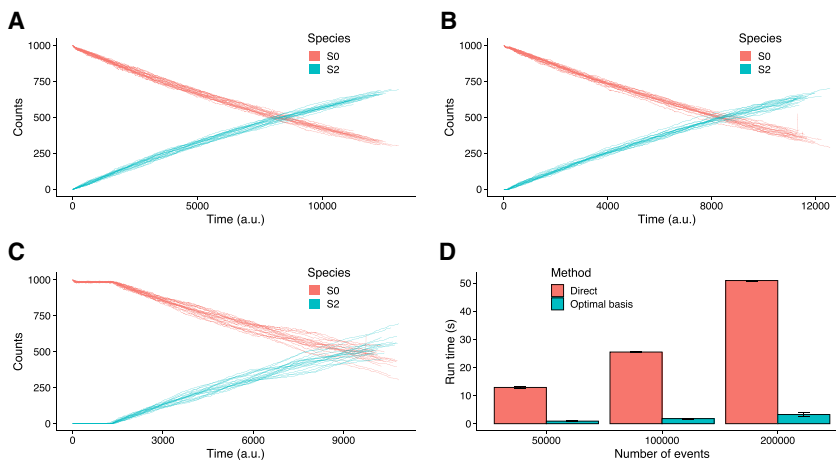


FIGURE 3 Comparison of the SSA-FBA simulation methods. Shown are trajectories from three sets of 20 representative simulations, each consisting of 3000 reaction execution events: (A) direct SSA-FBA simulation of toy model (equivalent to $\Delta_{event} = 0$); (B) approximate SSA-FBA simulation of toy model with $\Delta_{event} = 50$; and (C) approximate SSA-FBA simulation of toy model with $\Delta_{event} = 500$. (D) Comparison of run times for direct and optimal basis simulation methods. Error bars are SDs for run times across four replicate simulations involving 50,000, 100,000, or 200,000 FBA-bound updating events (described in Supporting materials and methods, Section S2). To see this figure in color, go online.

time by an order of magnitude over the latter. Intuitively, the performance of the optimal basis algorithm relative to the direct method will depend on the level of connectivity between the embedded FBA problem and remaining SSA component of a model. Greatest performance increases will be seen in models in which most reaction execution events do not require updating many FBA bounds so that the computational cost associated with validating the current optimal basis is substantially less than that associated with solving the embedded LP problem. Given the typical ways of representing biological coupling between metabolic and macromolecular portions of models (e.g., enzyme capacity constraints in which most enzymes are involved in the catalysis of a single reaction (20,40)), it follows that the connectivity of biologically realistic reaction networks implies the optimal basis algorithm will significantly improve simulation performance for the vast majority of cases.

Case study: reduced single-cell model of *M. pneumoniae*

In this section, we present results from a case study using SSA-FBA to simulate the dynamics of metabolism in a reduced model of a single *M. pneumoniae* cell to understand how variability in single-cell metabolism is driven by the intrinsic stochasticity of gene expression. *Mycoplasma* have the smallest genomes among known freely living cells, and by the law of large numbers, their small sizes therefore mean they should exhibit the largest effects of stochasticity in the absence of regulation: the relative absence of genetically encoded regulation compared with higher organisms means that stochastic effects are likely to play a larger role in their metabolic behavior. Furthermore, *M. pneumoniae* is one of the best-characterized members of this family (46), which makes it an excellent case study for understanding the behavioral effects of stochasticity in single-cell metabolism. Its close relative *M. genitalium* was previously used to build a whole-cell model (28), to which SSA-FBA could also be applied in principle along with genome-scale metabolic reaction networks from other organisms. The purpose of this section is not to study an entire whole-cell model, however, because the complexity involved in doing so would extend far beyond the main scope of this article, the goal of which is to introduce SSA-FBA as a modeling framework and demonstrate its feasibility by simulating a model of reasonable size and biological accuracy.

We constructed a model of *M. pneumoniae* consisting of 505 biochemical reactions that account for the physiology of metabolism and the intrinsic drivers of stochasticity in gene transcription, translation, and macromolecular degradation. Most parameter values were derived from metabolomics, transcriptomics, and proteomics data about *M. pneumoniae* (47–50), with the exception of rate con-

stants for metabolic reactions (for which many fewer experimental observations are available) that are drawn from a variety of other bacteria (51). The model contains a metabolic reaction network with 86 metabolic reactions (Fig. 4 A), including a biosynthetic pseudoreaction that constitutes the objective function of the embedded FBA problem. The biosynthetic pseudoreaction was constructed to be internally consistent with the SSA-only reactions that model gene expression as described below and does not account for additional maintenance energy costs not included in the model—a complete description of the biosynthetic pseudoreaction is provided in the relevant subsection of [Supporting materials and methods](#), Section S3. There are 81 protein species that either function individually or in complex (27 complexes in total) as enzymes or transporters regulating flux through the FBA reactions or serve a direct role in gene expression (ribosomal proteins, RNA polymerase, or RNase subunits). Furthermore, each protein species is associated with a messenger RNA molecule (in addition to the three ribosomal RNAs) and four reactions corresponding to RNA transcription, RNA degradation, protein synthesis, and protein degradation, which together with complexation reactions, makes 420 SSA-only reactions.

As described previously, it is natural that metabolite species (NTPs and amino acids (AAs)) appearing as reactants in the rate laws for the SSA-only reactions governing macromolecular synthesis and degradation are viewed as external species of the metabolic reaction network along with extracellular substrates, whereas all other metabolites are considered internal. Therefore, reactions of the embedded FBA model involved in the production or consumption of NTPs and AAs or extracellular transport are naturally assigned to the SSA-FBA reaction subset. The remaining reactions of the embedded FBA model (the substrates or products of which do not involve NTPs, AAs, or extracellular substrates) are consequently assigned to the FBA-only subset. [Supporting materials and methods](#) contains all of the compartments, species, reactions, rate laws, and parameters; [Supporting materials and methods](#), Section S3 provides an extended description of the model.

The reduced single-cell model is based partly on metabolic rate constants measured for bacteria distantly related to *M. pneumoniae* and so should not be considered a precise description of *Mycoplasma* physiology. Indeed, even the most detailed constraint-based, genome-scale models attempting to provide descriptions of entire cell physiology (53) must contend with the problem of missing parameter values, and the dependence of their completeness on future experimental research is something that genome-scale SSA-FBA simulations are subject to also. The average fold variation for wild-type bacterial enzymatic reaction rate constants (k_{cat} -values) in BRENDA (51) is 3951.6, as measured by $(k_{cat}^{max} - k_{cat}^{min})/k_{cat}^{min}$ (where k_{cat}^{max} and k_{cat}^{min} are the largest and smallest, respectively, k_{cat} -values reported for that enzyme) and averaged across all enzymes included

in the model, suggesting one or more of these model parameters require calibration to, at the very least, generate predictions consistent with the observed growth physiology of *M. pneumoniae*.

Here, one advantage coming from computational efficiency of the optimal basis SSA-FBA simulation algorithm is making it possible to parsimoniously search metabolic parameter space and establish parameter sensitivities by individually varying enzymatic reaction rate constants through repeated simulation. As an illustrative example, we focus on the sensitivity of growth rate to variations in the k_{cat} -value of guanylate kinase (GMK; gene MPN246) as a key enzyme (a phosphotransferase) forming part of the phosphotransfer network responsible for the homeostasis of NTPs and cellular energetics (54). To generate each set of simulation results for a given GMK k_{cat} -value, we ran 1000 instances of the single-cell *M. pneumoniae* model using the optimal basis SSA-FBA method initialized with counts of species randomly sampled from a Poisson distribution parametrized by their mean values across the cell cycle. In each case, we analyzed simulations over a 25-min interval of biological time, across which cell volume is assumed to be approximately constant given the relatively long (6–8 h) doubling time of *M. pneumoniae* (46), excluding the first 25 min of each 50-min simulation to prevent initial transients from contaminating the analysis.

The k_{cat} -value of GMK was varied from its lowest measured value (equal to the k_{cat} of adenylate kinase (ADK) in *Bacillus subtilis*, a much faster growing gram-positive bacteria) by increasing in 10-fold increments toward its highest estimated value (ADK in *Vibrio natriegens*). Notably, we found that the distribution of SSA execution events per simulation becomes bimodal as the GMK reaction rate constant increases (Fig. 4 B). This bimodal distribution at larger GMK k_{cat} -values computationally separates simulations into two distinct groups with a low (less than 10,000 SSA execution events) or high (more than 10,000 SSA execution events) number of SSA execution events per simulation, which are biologically associated with significantly low or high time-averaged GTP concentrations: 0.04 ± 0.06 mM vs. 0.47 ± 0.29 mM (mean \pm standard deviation (SD)) in low versus high group, respectively (Fig. 4 C). Fig. 4 C also shows that for the largest value of the GMK reaction rate constant tested, simulations in the low GTP group (146 out of 1000 simulations) were associated with a significantly lower average proportion of SSA execution events corresponding to metabolic reactions compared with simulations in the high GTP group: 79.36 vs. 99.16% in low GTP group versus high GTP group, respectively, implying that higher time-averaged GTP concentrations are associated with increased metabolic activity. Conversely, a significantly relatively higher average proportion of SSA execution events correspond to translation reactions in the low GTP group compared with the high GTP group: 13.46% compared with 0.50% in low GTP group

versus high GTP group, respectively, consistent with the fact that translation is a major consumer of intracellular GTP. Given the critical importance of ATP as source of energy for intracellular reactions supporting growth, we then studied the distribution of time-averaged ATP concentrations across the population of simulated cells. We found that the time-averaged concentrations ATP across the group of metabolically more active cells is positively skewed in qualitative agreement with experimental measurements of ATP concentrations inside individual bacteria (52) (Fig. 4 D): 1.28 ± 0.75 mM vs. 1.54 ± 1.22 mM (mean \pm SD) and skew values of 1.16 vs. 2.20 for simulation versus experiment, respectively. Conversely, the distribution of ATP concentrations in metabolically less active cells was slightly negatively skewed: 0.81 ± 0.22 mM (mean \pm SD) and skew value of -0.35 .

To understand whether multiple optimal solutions of the embedded FBA problem explain the bimodal simulation behavior observed at larger GMK k_{cat} -values, we performed flux variation analysis (44) with the metabolic reaction network at steady state. Flux variation analysis revealed that the reactions catalyzed by GMK, ADK, and phosphoglycerate kinase (PGK; ATP and GTP production) are able to carry arbitrarily large flux values within the optimal solution space defined by maximizing flux through the biosynthetic pseudoreaction. Correspondingly, elementary flux mode enumeration (55) used to identify all minimal pathways through the metabolic reaction network revealed a single internal cycle (and its reverse) that contains the support of the GMK, ADK, and PGK reactions without net metabolite production or consumption. Such internal cycles, not involving the primary exchange reactions representing the exchange of material between the model and the environment, are known to violate the first law of thermodynamics in conventional constraint-based models (43), but in the reduced *M. pneumoniae* model, GMK, ADK, and PGK catalyze SSA-FBA reactions involved in internal exchange of NTPs. We confirmed that simulations within the group of metabolically more active cells have significantly higher mean numbers of these four reactions executing across the course of a simulation than any other SSA-FBA reaction ($p < 2.2 \times 10^{-16}$, based on one sample *t*-test between the mean number of execution times of a given SSA-FBA reaction compared with that of all others), which is a result of the internal loop being able to carry a larger flux value in the optimal solution space when the GMK reaction rate is bounded by a higher k_{cat} -value. We also found a significant positive correlation between time-averaged ATP concentrations and ADK but not PGK enzyme levels ($R = 0.18$, $p < 10^{-7}$, and $R = 0.04$, $p = 0.24$, respectively, based on Pearson's product moment correlation coefficient), which were positively skewed (1.24 ± 0.26 μ M and 12.35 ± 1.52 μ M (mean \pm SD) and skew values 1.00 and 1.67, respectively (Fig. 4 E), suggesting that asymmetric metabolite distributions are perhaps the result of fluctuating ADK

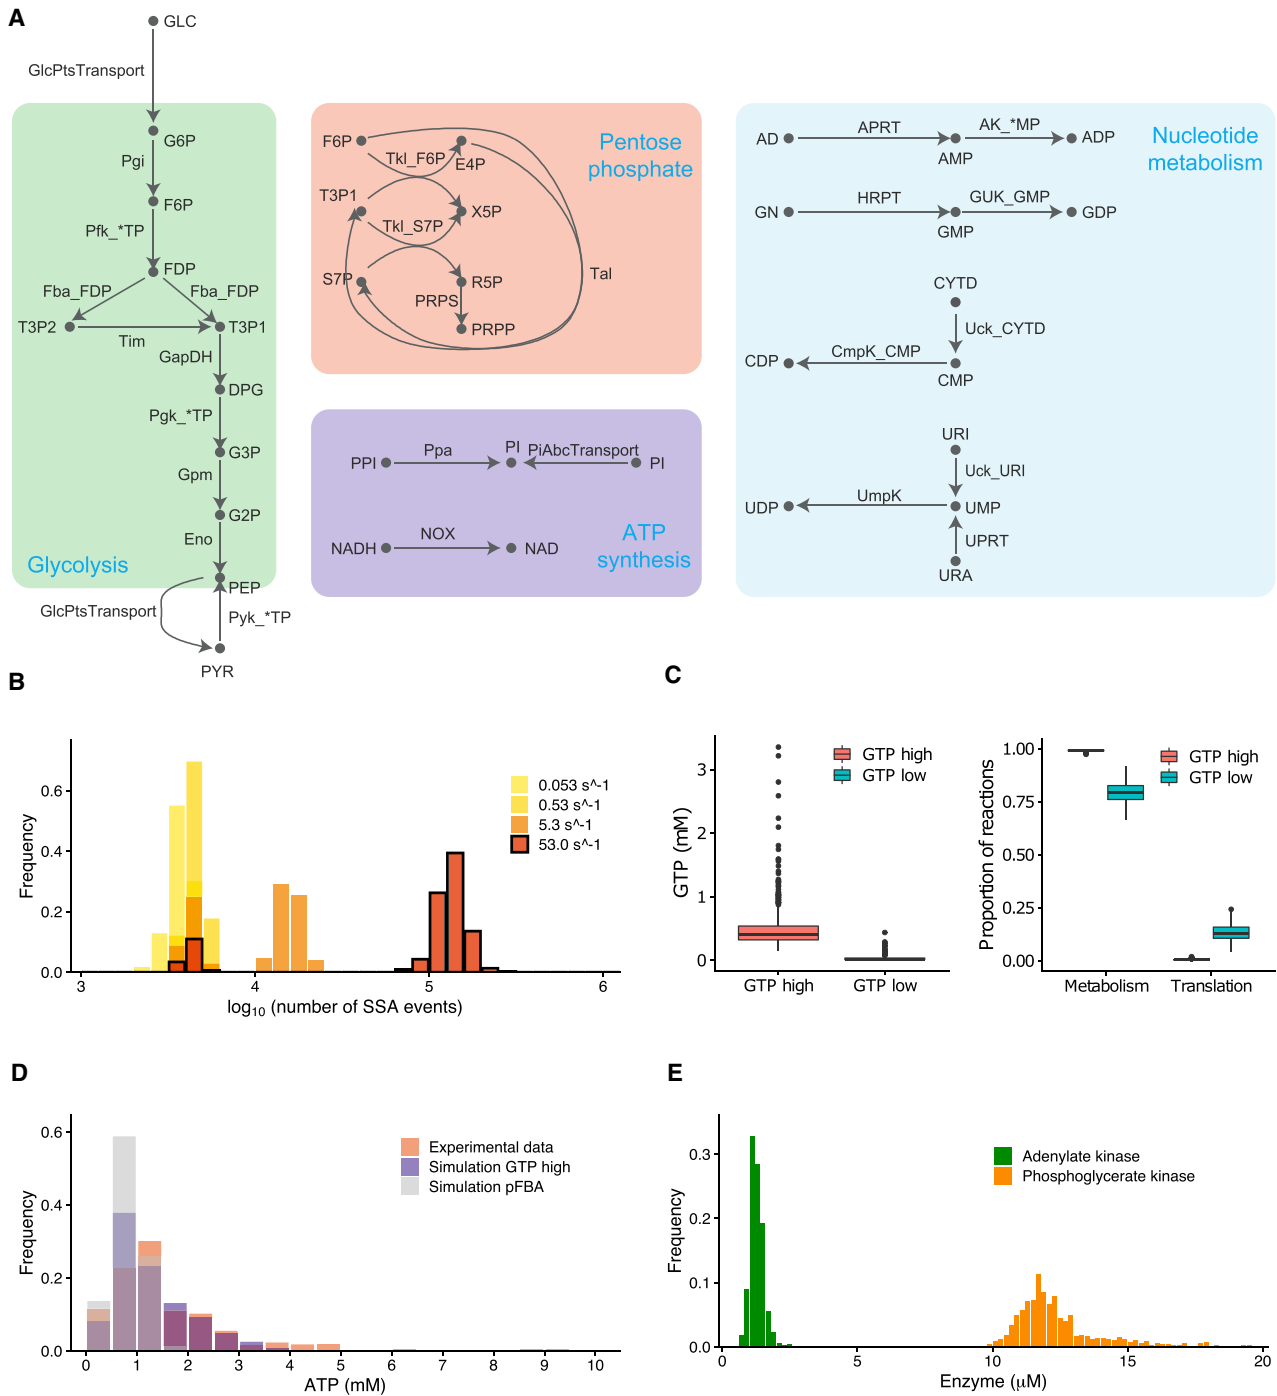


FIGURE 4 Single-cell simulation of *M. pneumoniae* metabolism using SSA-FBA. (A) Metabolic reaction network used in the model annotated with species listed in Supporting materials and methods. (B) Comparison of distributions of the number of SSA execution events per simulation for different values of the GMK reaction rate constant. For reference, the bimodal distribution corresponding to $k_{cat} = 53.0 \text{ s}^{-1}$ is outlined in black. (C) Box plots showing significant differences between GTP concentrations ($p < 2.2 \times 10^{-16}$) and the relative proportions of reaction execution events that correspond to metabolic reactions ($p < 2.2 \times 10^{-16}$) and translation reactions ($p < 2.2 \times 10^{-16}$), respectively, between the low and high GTP groups. (D) Distributions of time-averaged ATP concentrations from simulations with the GTP-high group and pFBA compared with the experimental distribution from (52). (E) Distribution of time-averaged ADK and PGK concentrations from simulations with the GTP-high group. To see this figure in color, go online.

levels controlling the relative levels of flux through the internal loop. Conversely, in the group of metabolically less active cells, reactions catalyzed by GTP phosphofrukti-

nase, PGK, and GTP pyruvate kinase along with the glucose import reaction are significantly overrepresented in simulations ($p < 2.2 \times 10^{-16}$), without significant correlation

between time-averaged ATP concentrations and ADK enzyme levels.

Because the internal cycle involving ADK responsible for positively skewed time-averaged ATP concentrations is deemed thermodynamically infeasible by constraint-based modeling criteria (43), we next employed lexicographic optimization with the biosynthetic pseudoreaction and a second parsimonious FBA (pFBA) objective (minimizing the sum of absolute flux values; see (56) and [Supporting materials and methods](#), Section S3). Although imposing a strong assumption on the biology of *M. pneumoniae* metabolism, pFBA is guaranteed to return a set of optimal SSA-FBA propensity values that are calculated without the involvement of internal cycles. Subsequently, the distribution of SSA execution events per simulation with pFBA retained a single mode even at the highest value of the GMK reaction rate constant and SSA-FBA execution events associated with ATP phosphofructokinase, PGK, ATP pyruvate kinase, and glucose import were significantly overrepresented ($p < 2.2 \times 10^{-16}$). Mean ATP concentrations with the pFBA objective were closer to those of previously metabolically less active cells, although the distribution retained a slight positive skew: 0.83 ± 0.01 mM (mean \pm SD) and a skew value of 0.30 (Fig. 4 D). Taken together, the requirement to include thermodynamically infeasible internal cycles to reproduce experimentally observed ATP distributions is a critical shortcoming of the reduced *M. pneumoniae* single-cell model that emerges because SSA-FBA reactions involved in internal NTP exchange are coupled to macromolecular synthesis and degradation. This issue would not arise in typical population-based DFBA models in which the embedded FBA problem is only coupled to the ODE via its primary exchange reactions. It implies that additional nongrowth-associated NTP maintenance reactions should be incorporated into single-cell models to account for alternative NTP requirements (57) and more formally mediate internal NTP exchange, possibly with additional constraints that impose conservation on certain metabolite pools (58).

DISCUSSION

In this article, we have presented SSA-FBA, the first framework, to our knowledge, with a realistic potential for simulating the complete dynamics of metabolic networks in single cells at the resolution of individual species and reactions. SSA-FBA is a hybrid method for embedding FBA in SSA that is well suited to simulating metabolic networks for which detailed kinetic information is often lacking. We also have developed an advanced optimal basis algorithm for efficiently executing SSA-FBA without approximation. A case study using our algorithm to simulate a reduced model for the metabolism of an individual *M. pneumoniae* cell demonstrates that SSA-FBA has the potential to reveal how stochasticity at the single-cell level contributes to meta-

bolic heterogeneity at the population level. Our results may help to identify the sources of variation observed in experimental measurements of important metabolites in single cells, indicating that reaction network coupling could contribute to population heterogeneity by amplifying or attenuating the sources noise in an integrative fashion. SSA-FBA can therefore be used to probe metabolic heterogeneity in a manner complementary to alternative existing experimental and computational methods (15,22,25,27).

Limitations of this study include various assumptions involved in the separation of timescales between reactions from metabolic portions of models and reactions responsible for macromolecular synthesis and degradation. In particular, although the faster timescales of metabolic reactions and higher abundance levels of metabolites have been widely used to justify a reduction of the CME or deterministic descriptions of metabolism (25,37–39), the validity of these assumptions should be questioned for each individual representation of single-cell metabolism before implementation of SSA-FBA. We have also highlighted two additional restrictions on SSA-FBA that reflect general limitations of any metabolic modeling framework based on an LP formulation but can be overcome using extensions of the work presented here: first, uniqueness of SSA-FBA propensity values is not guaranteed, although this can be achieved by implementing lexicographic optimization compatible with the efficient SSA-FBA simulation algorithm; second, although not possible to simulate using the efficient algorithm, SSA-FBA models can be generalized to include scenarios where a nonlinear optimization problem is used to represent metabolism. Finally, the general lack of experimentally determined parameter values that plagues current whole-cell modeling attempts (28,29,53) serves as a current limitation to building large-scale, single-cell models of metabolism, but as further advancements are made in data collection, SSA-FBA will retain an advantage over alternative frameworks that depend on precise knowledge of kinetic parameters.

Future extensions of our work will involve applying SSA-FBA to larger, more realistic models of entire cells such as those based on resource balance analysis (53), in addition to reconsidering how standard constraint-based formulations of metabolite pools, energy maintenance, and objective functions should be adapted to suit the biological nature of single-cell biology. More complex models of single-cell metabolism could include mechanisms that impart regulatory control of stochasticity (59) and pave the way for combining insights from simulation with experimental advances in microfluidics (60) or real-time quantification of RNA translation events within individual cells (61). From an algorithmic perspective, SSA-FBA could be extended to incorporate additional timescales governed by continuous stochastic or deterministic processes (e.g., (62)) and its computational efficiency perhaps further enhanced through parallelization methods (63).

SUPPORTING MATERIAL

Supporting material can be found online at <https://doi.org/10.1016/j.bpj.2021.10.038>.

AUTHOR CONTRIBUTIONS

D.S.T. designed the research and implemented algorithms and software with conceptual feedback from the other coauthors. J.R.K. designed the case study, and D.S.T. implemented and performed the simulations. D.S.T. wrote the article with input from all authors.

ACKNOWLEDGMENTS

We thank M. E. Beber and J. Carrasco Muriel for their close collaboration with the first author on a related project that provided some of the motivation for this study and Maria Lluch-Senar, Veronica Llorens, and Luis Serrano for discussions on the physiology of mycoplasmas.

Part of this work was completed while D.S.T. was a Simons Foundation Fellow of the Life Sciences Research Foundation, hosted by Columbia University Irving Medical Center. The work of J.R.K. was supported by National Science Foundation award 1649014 and National Institutes of Health award R35GM119771.

REFERENCES

- Schier, A. F. 2020. Single-cell biology: beyond the sum of its parts. *Nat. Methods*. 17:17–20.
- Angerer, P., L. Simon, ..., F. J. Theis. 2017. Single cells make big data: new challenges and opportunities in transcriptomics. *Curr. Opin. Syst. Biol.* 4:85–91.
- Stuart, T., and R. Satija. 2019. Integrative single-cell analysis. *Nat. Rev. Genet.* 20:257–272.
- Ackermann, M. 2015. A functional perspective on phenotypic heterogeneity in microorganisms. *Nat. Rev. Microbiol.* 13:497–508.
- Schreiber, F., and M. Ackermann. 2020. Environmental drivers of metabolic heterogeneity in clonal microbial populations. *Curr. Opin. Biotechnol.* 62:202–211.
- Dagogo-Jack, I., and A. T. Shaw. 2018. Tumour heterogeneity and resistance to cancer therapies. *Nat. Rev. Clin. Oncol.* 15:81–94.
- Xiao, Z., J. W. Locasale, and Z. Dai. 2020. Metabolism in the tumor microenvironment: insights from single-cell analysis. *OncolImmunology*. 9:1726556.
- Rubakhin, S. S., E. J. Lanni, and J. V. Sweedler. 2013. Progress toward single cell metabolomics. *Curr. Opin. Biotechnol.* 24:95–104.
- Fessenden, M. 2016. Metabolomics: small molecules, single cells. *Nature*. 540:153–155.
- Emara, S., S. Amer, ..., T. Masujima. 2017. Single-cell metabolomics. In *Metabolomics: From Fundamentals to Clinical Applications* A. Sussulini, ed., Springer, pp. 323–343.
- Duncan, K. D., J. Fyrestam, and I. Lanekoff. 2019. Advances in mass spectrometry based single-cell metabolomics. *Analyst (Lond.)*. 144:782–793.
- Damiani, C., D. Gaglio, ..., M. Vanoni. 2020. Systems metabolomics: from metabolomic snapshots to design principles. *Curr. Opin. Biotechnol.* 63:190–199.
- Hanahan, D., and R. A. Weinberg. 2011. Hallmarks of cancer: the next generation. *Cell*. 144:646–674.
- Pavlova, N. N., and C. B. Thompson. 2016. The emerging hallmarks of cancer metabolism. *Cell Metab.* 23:27–47.
- Evers, T. M. J., M. Hochane, ..., A. Mashaghi. 2019. Deciphering metabolic heterogeneity by single-cell analysis. *Anal. Chem.* 91:13314–13323.
- Edwards, J. S., M. Covert, and B. Palsson. 2002. Metabolic modelling of microbes: the flux-balance approach. *Environ. Microbiol.* 4:133–140.
- O'Brien, E. J., J. M. Monk, and B. O. Palsson. 2015. Using genome-scale models to predict biological capabilities. *Cell*. 161:971–987.
- Angione, C. 2019. Human systems biology and metabolic modelling: a review- from disease metabolism to precision medicine. *BioMed Res. Int.* 2019:8304260.
- Yang, L., A. Ebrahim, ..., B. O. Palsson. 2019. DynamicME: dynamic simulation and refinement of integrated models of metabolism and protein expression. *BMC Syst. Biol.* 13:2.
- Waldherr, S., D. A. Oyarzún, and A. Bockmayr. 2015. Dynamic optimization of metabolic networks coupled with gene expression. *J. Theor. Biol.* 365:469–485.
- Kiviet, D. J., P. Nghe, ..., S. J. Tans. 2014. Stochasticity of metabolism and growth at the single-cell level. *Nature*. 514:376–379.
- Wehrens, M., F. Büke, ..., S. J. Tans. 2018. Stochasticity in cellular metabolism and growth: approaches and consequences. *Curr. Opin. Syst. Biol.* 8:131–136.
- Damiani, C., D. Maspero, ..., G. Mauri. 2019. Integration of single-cell RNA-seq data into population models to characterize cancer metabolism. *PLoS Comput. Biol.* 15:e1006733.
- Levine, E., and T. Hwa. 2007. Stochastic fluctuations in metabolic pathways. *Proc. Natl. Acad. Sci. USA*. 104:9224–9229.
- Tonn, M. K., P. Thomas, ..., D. A. Oyarzún. 2019. Stochastic modeling reveals mechanisms of metabolic heterogeneity. *Commun. Biol.* 2:108.
- Thomas, P., G. Terradot, ..., A. Y. Weiße. 2018. Sources, propagation and consequences of stochasticity in cellular growth. *Nat. Commun.* 9:4528.
- Tonn, M. K., P. Thomas, ..., D. A. Oyarzún. 2020. Computation of single-cell metabolite distributions using mixture models. *Front. Cell Dev. Biol.* 8:614832.
- Karr, J. R., J. C. Sanghvi, ..., M. W. Covert. 2012. A whole-cell computational model predicts phenotype from genotype. *Cell*. 150:389–401.
- Karr, J. R., K. Takahashi, and A. Funahashi. 2015. The principles of whole-cell modeling. *Curr. Opin. Microbiol.* 27:18–24.
- Gillespie, D. T. 1976. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.* 22:403–434.
- Gillespie, D. T. 1977. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* 81:2340–2361.
- Mahadevan, R., J. S. Edwards, and F. J. Doyle, III. 2002. Dynamic flux balance analysis of diauxic growth in *Escherichia coli*. *Biophys. J.* 83:1331–1340.
- Schnoerr, D., G. Sanguinetti, and R. Grima. 2017. Approximation and inference methods for stochastic biochemical kinetics – a tutorial review. *J. Phys. A Math. Theor.* 50:093001.
- Elowitz, M. B., A. J. Levine, ..., P. S. Swain. 2002. Stochastic gene expression in a single cell. *Science*. 297:1183–1186.
- Gillespie, D. T. 1992. A rigorous derivation of the chemical master equation. *Physica A*. 188:404–425.
- Rao, C. V., and A. P. Arkin. 2003. Stochastic chemical kinetics and the quasi-steady-state assumption: application to the Gillespie algorithm. *J. Chem. Phys.* 118:4999–5010.
- Cao, Y., D. T. Gillespie, and L. R. Petzold. 2005. The slow-scale stochastic simulation algorithm. *J. Chem. Phys.* 122:14116.
- Cao, Y., D. T. Gillespie, and L. R. Petzold. 2005. Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems. *J. Comput. Phys.* 206:395–411.

39. Smith, S., C. Cianci, and R. Grima. 2015. Model reduction for stochastic chemical systems with abundant species. *J. Chem. Phys.* 143:214105.
40. Sánchez, B. J., C. Zhang, ..., J. Nielsen. 2017. Improving the phenotype predictions of a yeast genome-scale metabolic model by incorporating enzymatic constraints. *Mol. Syst. Biol.* 13:935.
41. Gibson, M. A., and J. Bruck. 2000. Efficient exact stochastic simulation of chemical systems with many species and many channels. *J. Phys. Chem. A.* 104:1876–1889.
42. Gillespie, D. T. 2001. Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.* 115:1716–1733.
43. Henry, C. S., L. J. Broadbelt, and V. Hatzimanikatis. 2007. Thermodynamics-based metabolic flux analysis. *Biophys. J.* 92:1792–1805.
44. Mahadevan, R., and C. H. Schilling. 2003. The effects of alternate optimal solutions in constraint-based genome-scale metabolic models. *Metab. Eng.* 5:264–276.
45. Harwood, S. M., K. Höffner, and P. I. Barton. 2016. Efficient solution of ordinary differential equations with a parametric lexicographic linear program embedded. *Numer. Math.* 133:623–653.
46. Waites, K. B., and D. F. Talkington. 2004. *Mycoplasma pneumoniae* and its role as a human pathogen. *Clin. Microbiol. Rev.* 17:697–728.
47. Güell, M., V. van Noort, ..., L. Serrano. 2009. Transcriptome complexity in a genome-reduced bacterium. *Science.* 326:1268–1271.
48. Kühner, S., V. van Noort, ..., A.-C. Gavin. 2009. Proteome organization in a genome-reduced bacterium. *Science.* 326:1235–1240.
49. Maier, T., A. Schmidt, ..., L. Serrano. 2011. Quantification of mRNA and protein and integration with protein turnover in a bacterium. *Mol. Syst. Biol.* 7:511.
50. Yus, E., V. Lloréns-Rico, ..., L. Serrano. 2019. Determination of the gene regulatory network of a genome-reduced bacterium highlights alternative regulation independent of transcription factors. *Cell Syst.* 9:143–158.e13.
51. Scheer, M., A. Grote, ..., D. Schomburg. 2011. BRENDA, the enzyme information system in 2011. *Nucleic Acids Res.* 39:D670–D676.
52. Yaginuma, H., S. Kawai, ..., H. Imamura. 2014. Diversity in ATP concentrations in a single bacterial cell population revealed by quantitative single-cell imaging. *Sci. Rep.* 4:6522.
53. Bulović, A., S. Fischer, ..., A. Goelzer. 2019. Automated generation of bacterial resource allocation models. *Metab. Eng.* 55:12–22.
54. Dzeja, P. P., and A. Terzić. 2003. Phosphotransfer networks and cellular energetics. *J. Exp. Biol.* 206:2039–2047.
55. Terzer, M., and J. Stelling. 2008. Large-scale computation of elementary flux modes with bit pattern trees. *Bioinformatics.* 24:2229–2235.
56. Desouki, A. A., F. Jarre, ..., M. J. Lercher. 2015. CycleFreeFlux: efficient removal of thermodynamically infeasible loops from flux distributions. *Bioinformatics.* 31:2159–2165.
57. Thiele, I., and B. O. Palsson. 2010. A protocol for generating a high-quality genome-scale metabolic reconstruction. *Nat. Protoc.* 5:93–121.
58. De Martino, A., D. De Martino, ..., A. Pagnani. 2014. Identifying all moiety conservation laws in genome-scale metabolic networks. *PLoS One.* 9:e100750.
59. Borri, A., P. Palumbo, and A. Singh. 2016. Impact of negative feedback in metabolic noise propagation. *IET Syst. Biol.* 10:179–186.
60. Leygeber, M., D. Lindemann, ..., D. Kohlheyer. 2019. Analyzing microbial population heterogeneity-expanding the toolbox of microfluidic single-cell cultivations. *J. Mol. Biol.* 431:4569–4588.
61. Morisaki, T., K. Lyon, ..., T. J. Stasevich. 2016. Real-time quantification of single RNA translation dynamics in living cells. *Science.* 352:1425–1429.
62. Haseltine, E. L., and J. B. Rawlings. 2002. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J. Chem. Phys.* 117:6959.
63. Goldberg, A. P., D. R. Jefferson, ..., J. R. Karr. 2020. Exact parallelization of the stochastic simulation algorithm for scalable simulation of large biochemical networks. *arXiv*, arXiv:2005.05295 <https://arxiv.org/abs/2005.05295>.