Finding Equilibria in Large Games using Variational Inequalities

Ian Gemp and **Sridhar Mahadevan**

School of Computer Science
University of Massachusetts
Amherst, MA 01003
(imgemp,mahadeva)@cs.umass.edu

Abstract

In this paper, we explore an approach to computational game theory based on variational inequalities (VIs). VIs represent a comprehensive framework that provides a way to model and analyze both cooperative and non-cooperative games. Given the potentially large size of real-world games, suitable algorithms must be designed that can scale gracefully with the dimension of the problems (e.g., number of players). In this paper, we explore the effectiveness of novel Runge-Kutta methods on finding equilibrium solutions to two real-world games defined by oligopolistic economies.

1 Introduction

A significant focus of game theory is the search for equilibria. Equilibrium solutions are important because they often represent the behavior of the game in steady state. If the equilibrium state is undesirable, we may try to change the game in some way to force a more desirable steady state behavior. On the other hand, if the steady state is already satisfactory, we may still change the game in order to accelerate convergence to the equilibrium. The identification, analysis, and manipulation of equilibria represents a computationally challenging problem.

In this paper, we explore an approach to computational game theory in AI, based on the framework of variational inequalities (VIs). Originally proposed in the context of solving partial differential equations in mechanics (Hartman and Stampacchia 1966), VIs gained popularity in the finite-dimensional setting when the traffic network equilibrium problem was formulated as a finite-dimensional VI (Dafermos 1980). This advance inspired much follow-on research, showing that a variety of equilibrium problems in economics, game theory, and manufacturing could also be formulated as finite-dimensional VIs – the books by Nagurney (Nagurney 1999; Nagurney and Zhang 1996) and Facchinei and Pang (Facchinei and J. 2003) provide a detailed introduction to the theory and applications of VI.

We apply VIs to the problem of modeling large economic games, where the challenging computational problem is to find equilibrium solutions that balance numerous conflicting objectives. In order to solve these large games, we will require fast, scalable algorithms suitable to our problems. The primary purpose of this paper is to explain how the theory of VIs provides valuable computational tools for solving large competitive systems as well as present a suitable algorithm for solving such systems.

Section 2 describes two real-world oligopolistic economies involving a sustainable freightage network and a service-oriented internet network. Section 3 provides a brief overview of VIs and describes standard algorithms for solving VIs as well as a novel and general Runge-Kutta algorithmic framework that we propose for large domains. In Section 4, we explain the Runge-Kutta (RK) family of methods along with their associated adaptive stepsize scheme. Section 5 compares RK on VI formulations of a sustainable freight supply chain network and a next-generation economic model of the Internet. Experiments in both these domains show significant benefits of our proposed novel RK method.

2 Domain Backgrounds

We begin by describing two sustainable network domains used to empirically validate our approach to solving large games.

2.1 Sustainable Freightage Network

The Commission for Environmental Cooperation released a report in 2010 focusing on reducing the greenhouse gas (GHG) emissions from freight transportation in North America[1]. This report revealed that while "light-duty vehicle GHG emissions are projected...to decline nearly 12%..., Freight trucks, on the other hand, show a projected 20% increase in emissions." Among the commission's key findings is the need for the "greening" of supply-chain management. While some changes to the supply chain like reduced fuel consumption clearly reduce business costs, others may help to mitigate "reputational risk". An important objective for the next generation of supply-chain models is to incorporate these factors in order to reflect the changing goals. The network diagram associated with this problem is shown below

---

[1]http://www3.cec.org/islandora/en/item/4237-destination-sustainability-reducing-greenhouse-gas-emissions-from-freight-en.pdf

in Figure 1, which is based on a formulation proposed in (Nagurney, Yu, and Floden 2013).


Figure 1: A "green" economic model of the supply chain proposed in (Nagurney, Yu, and Floden 2013). Firms are modeled as playing a Cournot-Nash game, competing on the basis of product flow and frequency of operation. Demand markets consisting of individuals or groups of users choose between the various products offered by the firms.

In this network model, I firms manufacture products which are then either transported directly to retailers (demand markets) or to storage facilities for later distribution. The products in this economy are substitutable and distinguishable only by brand (e.g., milk). In addition, we assume knowledge of the demand functions stating the prices markets are willing to pay for quantities of each product. In Figure 1, the nodes from the top tier to the bottom tier represent in order the firms (i), manufacturing plants (M_m^i), storage warehouses ($D_{d,1}^i \& D_{d,2}^i$), and demand markets (R_r). Each link in the network represents a process acting on the product between the origin and destination nodes. From the top tier to the bottom tier, the links represent manufacturing, transportation, storage, and distribution. Note that each $D_{d,1}^i$ and $D_{d,2}^i$ pair actually represents the same distribution center. This is because storage is a process that starts and ends in the same warehouse, hence the duplication of the nodes.

Each firm must decide how to optimally deliver its product to consumers given the allowable paths from its firm to the multiple demand markets. They do this by controlling their product flows (e.g., gallons of milk per day) and frequencies of operation (e.g., shipments per day) along paths in the network subject to capacity constraints (e.g., gallons of milk per shipment). For example, firm 1 may decide on two paths to optimize its supply chain: each day, ten 5-gallon shipments are manufactured at plant 1 and transported using mode 4 (barge) directly to retail market 1 and six 3-gallon shipments are manufactured at plant 1 as well but are then transported using mode 3 (truck) to warehouse 2 for storage until they are finally distributed to retail market 11.

The firms in the network continuously adjust their product flows and operation frequencies, optimizing their utilities, until any unilateral adjustment attempted by one firm is inherently detrimental to that firm's utility function. Rationally competing on the basis of product output is known as Cournot competition and the stalemate described is known as a Nash equilibrium hence this state is known as a Cournot-Nash equilibrium.

Given each firm's utility function and capacity constraints, we aim to find the corresponding steady-state product flows and frequencies of operation.

2.2 "Next Generation" Internet

Our second example is a "next generation" economic model of the Internet. Since its inception, the size of the Internet has exploded, greatly outpacing any economic regulations that may have been instantiated to manage its evolution. Proponents of net-neutrality have argued that communications service providers (CSPs) have enjoyed a sheltered economic sphere due to lack of suitable laws and, until recently, poor understanding of the underlying economics[2]. The communications industry is expected to soon enter a highly competitive global dynamic marketplace, one more capable of supporting the extremely diverse demand markets of today[3]. These expectations require a more general, abstract perspective of the internet in order to understand its future.

The game-theoretic model of this "next generation" Internet, a service-oriented Internet, has members of the network (see Figure 2) compete to maximize profits by adjusting the quantity, quality, and price of services delivered. Service providers (e.g., Netflix, Amazon) play a *Cournot-Nash* game controlling the quantities of services provided while network providers (e.g., Verizon, AT&T) play a *Bertrand game* controlling the delivery price as well as service quality. Consumers influence the network through demand functions dictating the prices they are willing to pay for specific quantities and qualities of services rendered.

3 Variational Inequalities

The two networks shown in Figure 1 and Figure 2 pose a challenging computational problem, and our proposed solution builds on the mathematical framework of variational inequalities (VIs). As many readers may be unfamiliar with the mathematics of VIs, we begin with a brief review.

3.1 Theory

The formal definition of a VI is as follows:

Definition 1. *The finite-dimensional variational inequality problem VI(F,K) involves finding a vector $x^* \in K \subset \mathbb{R}^n$ such that*

$$\langle F(x^*), x - x^* \rangle \geq 0, \ \forall x \in K$$

where $F : K \to \mathbb{R}^n$ is a given continuous function, K is a given closed convex set, and $\langle ., . \rangle$ is the standard inner product in $\mathbb{R}^n$.

Figure 3 provides a geometric interpretation of a variational inequality. The following general result characterizes when solutions to VIs exist:


Figure 2: A next-generation economic model of the Internet proposed in (Nagurney and Wolf 2014). Service providers (e.g., "Netflix", "Amazon", "IBM Cloud") are modeled as playing a Cournot-Nash game, competing on the basis of quantity. Network providers (e.g., "Verizon", "AT & T") are modeled as playing a Bertrand game, competing on the basis of prices. Demand markets consisting of individual users or groups of users choose between combinations of service providers and transport providers.

Theorem 1. *Suppose K is compact and that $F : K \to \mathbb{R}^n$ is continuous. Then, there exists a solution to VI(F, K).*

As Figure 3 shows, x^* is a solution to $VI(F, K)$ if and only if the angle between the vectors $F(x^*)$ and $x - x^*$, for any vector $x \in K$, is less than or equal to 90^0.


Figure 3: This figure provides a geometric interpretation of the variational inequality $VI(F, K)$. The mapping F defines a vector field over the feasible set K such that at the solution point x^*, the vector field $F(x^*)$ is directed inwards at the boundary, and $-F(x^*)$ is an element of the normal cone $C(x^*)$ of K at x^* where the normal cone $C(x^*)$ at the vector x^* of a convex set K is defined as $C(x^*) = \{y \in \mathbb{R}^n | \langle y, x - x^* \rangle \leq 0, \forall x \in K\}$.

The VI framework provides a mathematically elegant approach to model equilibrium problems in game theory (Fudenberg and Levine 1998; Nisan et al. 2007). A *Nash game* consists of m players, where player i chooses a strategy x_i belonging to a closed convex set $X_i \subset \mathbb{R}^n$. After executing the joint action, each player is penalized (or rewarded)

by the amount $F_i(x_1, \ldots, x_m)$, where $F_i : \mathbb{R}^n \to \mathbb{R}$ is a continuously differentiable function. A set of strategies $x^* = (x_1^*, \ldots, x_m^*) \in \Pi_{i=1}^M X_i$ is said to be in equilibrium if no player can reduce the incurred penalty (or increase the incurred reward) by unilaterally deviating from the chosen strategy. If each F_i is convex on the set X_i, then the set of strategies x^* is in equilibrium if and only if $\langle (x_i - x_i^*), \nabla_i F_i(x_i^*) \rangle \geq 0$. In other words, x^* needs to be a solution of the VI $\langle (x - x^*), f(x^*) \rangle \geq 0$, where $f(x) = (\nabla F_1(x), \ldots, \nabla F_m(x))$. Nash games are closely related to *saddle point* problems (Juditsky, Nemirovski, and others 2011a; 2011b; Liu, Mahadevan, and Liu 2012) where we are given a function $F : X \times Y \to \mathbb{R}$, and the objective is to find a solution $(x^*, y^*) \in X \times Y$ such that

$$F(x^*, y) \leq F(x^*, y^*) \leq F(x, y^*), \forall x \in X, \forall y \in Y. \quad (1)$$

Here, F is convex in x for each fixed y, and concave in y for each fixed x. Many equilibria problems in economics can be modeled using VIs (Nagurney 1998).

The algorithmic development of methods for solving VIs begins with noticing their connection to fixed point problems.

Theorem 2. *The vector x^* is the solution of VI(F,K) if and only if, for any $\alpha > 0$, x^* is also a fixed point of the map $x^* = P_K(x^* - \alpha F(x^*))$, where P_K is the projector onto convex set K.*

In terms of the geometric picture of a VI illustrated in Figure 3, this property means that the solution of a VI occurs at a vector x^* where the vector field $F(x^*)$ induced by F on K is normal to the boundary of K and directed inwards, so that the projection of $x^* - \alpha F(x^*)$ is the vector x^* itself. This property forms the basis for the projection class of methods that solve for the fixed point.

Definition 2. *A gap function is a function $\psi : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ which satisfies $\psi(X) \geq 0$ for all $X \in K$ and $\psi(X^*) = 0, X^* \in K$ if and only if X^* solves $VI(F, K)$.*

While algorithmic convergence is often judged by the difference between successive states ($|x_{k+1} - x_k| < \epsilon$), this crude approach is unsatisfactory in the presence of adaptive step sizes (e.g. a small step size can satisfy this criterium regardless of the iterates). Gap functions provide a superior alternative as convergence criteria in the presence of adaptive step size schemes. Numerous gap functions have been developed satisfying the properties above (Dutta 2012). One such function, useful in unbounded domains, was developed separately by Fukushima and Auchmuty, $g_\alpha(x)$. We will use $g_\alpha(x)$ later in our experiments to judge convergence.

$$g_\alpha(x) = \sup_{y \in K} \{ \langle F(x), x - y \rangle - \frac{\alpha}{2} ||x - y||^2 \} \quad (2)$$

3.2 Algorithms

The basic projection-based method (Algorithm 1) for solving VIs is based on Theorem 2 introduced earlier. Here, P_K is the orthogonal projector onto the convex set K. It can be shown that the basic projection algorithm solves any $VI(F, K)$ for which the mapping F is strongly monotone and Lipschitz smooth. A simple strategy is to set $D = I$

Algorithm 1 The Basic Projection Algorithm.

INPUT: Given VI(F,K), and a symmetric positive definite matrix D.

1: Set $k = 0$ and $x_k \in K$.
2: **repeat**
3:     Set $x_{k+1} \leftarrow P_K(x_k - \alpha D^{-1}F(x_k))$.
4:     Set $k \leftarrow k + 1$.
5: **until** $x_k = P_K(x_k - \alpha D^{-1}F(x_k))$.
6: Return x_k

---

where $\alpha < \frac{L^2}{2\mu}$, L is the Lipschitz smoothness constant, and μ is the strong monotonicity constant. Setting D equal to a constant in this manner recovers what is known as Euler's method and is the most basic algorithm for solving VIs.

The basic projection-based algorithm has two critical limitations. First, it requires that the mapping F be strongly monotone. If, for example, F is the gradient map of a continuously differentiable function, strong monotonicity implies the function must be strongly convex. Second, setting the parameter α requires knowing the Lipschitz smoothness L and the strong monotonicity parameter μ.

Algorithm 2 The Extragradient Algorithm.

INPUT: Given VI(F,K), and a scalar α.

1: Set $k = 0$ and $x_k \in K$.
2: **repeat**
3:     Set $y_k \leftarrow P_K(x_k - \alpha F(x_k))$.
4:     Set $x_{k+1} \leftarrow P_K(x_k - \alpha F(y_k))$.
5:     Set $k \leftarrow k + 1$.
6: **until** $x_k = P_K(x_k - \alpha F(x_k))$.
7: Return x_k

The extragradient method of Korpolevich (Korpelevich 1977) addresses some of these concerns, and is defined as Algorithm 2. The extragradient algorithm has been the topic of much attention in optimization since it was proposed, e.g., see (Peng and Yao 2008; Nesterov 2007).

The family of Runge-Kutta (RK) methods induced by Nagurney's general iterative scheme (Nagurney and Zhang 1996) is defined as Algorithm 3. We will explain the motivation behind RK methods in the following section by observing their role in the solution of ordinary differential equations (ODEs). An RK method is defined by its values a and b which are typically presented as a Butcher table. Heun-Euler and Cash-Karp refer to two tables that we will use in our experiments.

4   Runge-Kutta Algorithms

In this section, we provide some intuition and explanation for the strengths of Runge-Kutta methods and their associated stepsize scheme.

4.1   The Runge-Kutta Method for ODEs

Runge-Kutta methods are highly popular methods for solving systems of coupled first-order differential equations of

Algorithm 3 The General Runge-Kutta Algorithm.

INPUT: Given VI(F,K), lower-triangular matrix $a \in \mathbb{R}^{s-1 \times s-1}$, vector $b \in \mathbb{R}^s$, and a sequence of scalars α_k.

1: Set $k = 0$ and $x_k \in K$.
2: **repeat**
3:     Set $k_1 \leftarrow \alpha_k F(x_k)$
4:     Set $k_2 \leftarrow \alpha_k F(P_K(x_k - a_{21}k_1))$
5:     Set $k_3 \leftarrow \alpha_k F(P_K(x_k - a_{31}k_1 - a_{32}k_2))$
     $\vdots$
6:     Set $k_s \leftarrow \alpha_k F(P_K(x_k - a_{s1}k_1 - \ldots - a_{s,s-1}k_{s-1}))$
7:     Set $x_{k+1} \leftarrow P_K(x_k - \sum_{i=1}^s b_i k_i)$
8:     Set $k \leftarrow k + 1$.
9: **until** $x_k = P_K(x_k - \sum_{i=1}^s b_i k_i)$.
10: Return x_k

---

the form:
$$\frac{dx}{dt} = f(x,t), \ \ x(t_0) = x_0 \qquad (3)$$
The simplest explicit RK method is Euler's method:
$$x_{k+1} = x_k + \alpha f(x_k, t_k), \ \ t_{k+1} = t_k + \alpha \qquad (4)$$
Euler's method, while simple, is not very accurate because it only uses the derivative of the function at the beginning of the interval. More accurate methods can be designed that advance x_k by a weighted mean of the derivatives of the function in a neighborhood of (x_k, t_k). Runge-Kutta methods are crafted such that the locations and corresponding weights at which the derivatives are computed induce an approximate Taylor series expansion of the algorithm that matches the infinite Taylor series expansion of x up to some order p, $\mathcal{O}(h^p)$. The general explicit Runge-Kutta scheme is given below where c_i and a_{ij} designate which locations to inspect and b_i defines the weights. Euler's method corresponds to $s = 1, b_1 = 1, c_1 = 1$.
$$x_{k+1} = x_k + \Sigma_{i=1}^s b_i k_i, \ \ t_{k+1} = t_k + \alpha, \qquad (5)$$
$$k_i = \alpha f(x_k + \Sigma_{j=1}^{i-1} a_{ij} k_j, t_k + c_i \alpha) \qquad (6)$$
Research into applying higher order methods has been conducted (Sen and Shanno 2008), however, use of Runge-Kutta methods and their convenient adaptive step size schemes is unexplored as far as we know.

4.2   Using Adaptive Stepsizes in RK Methods

We now describe an additional enhancement of Runge-Kutta methods that automatically tunes the stepsize. These adaptive methods are designed to produce an estimate of the local truncation error of a single Runge-Kutta step. This can be accomplished by computing and comparing steps with two methods during each iteration of descent, however, more efficient methods make use of the same Runge-Kutta matrix, but differing weights, b_i. This is done by simultaneously using two methods, one with order p and one with order $p - 1$. The lower-order step is given by
$$x_{k+1}^* = x_k + \sum_{i=1}^s b_i^* k_i, \qquad (7)$$

where the k_i are the same as for the higher-order method. Then the error is

$$\Delta_{k+1} = x_{k+1} - x^*_{k+1} = \sum_{i=1}^{s}(b_i - b^*_i)k_i, \qquad (8)$$

which is $\mathcal{O}(h^p)$. Stepsizes can be updated as $\alpha_{k+1} \leftarrow \alpha_k \left| \frac{\Delta_0}{\Delta_{k+1}} \right|^{1/p}$, where Δ_0 is the desired accuracy.

Classical gradient rules commonly enforce diminishing stepsizes. The scheme above, however, describes a stepsize that depends on the local behavior of f and may possibly grow with successive iterations. In fact, when used in practice, the stepsize increases as the solution nears the equilibrium to account for the diminishing value of the vector field.

5 Experiments

We now compare the proposed RK family of methods against standard VI algorithms on the domains discussed in Section 2.

5.1 Sustainable Freightage Network Experiment

Our first example focuses on an emissions-conscious competitive supply chain network (Nagurney, Yu, and Floden 2013). We assume the governing equilibrium is Cournot-Nash and the utility functions are all concave and fully differentiable. This establishes the equivalence between the equilibrium state we are searching for and the variational inequality to be solved where the F mapping is a vector consisting of the negative gradients of the augmented Lagrangian utility functions for each firm.

$\langle F(X^*), X - X^* \rangle \geq 0, \forall X \in \mathcal{K}$, where $X = (x, \gamma, \lambda) \in \mathbb{R}^{N_X}_+$
$F^1(X) = \frac{\partial C_p(x)}{\partial x_p} + \omega_i \frac{\partial E_p(x,\gamma)}{\partial x_p} + \sum_{a \in L^i} \lambda_a \delta_{ap} - \rho_{ik}(x)$
$- \sum_{l=1}^{n_R} \frac{\partial \rho_{il}(x)}{\partial x_p} \sum_{q \in P^i_l} x_q; p \in P^i_k; i = 1, \dots, I; k = 1, \dots, n_R$
$F^2(X) = \frac{\partial g_a(\gamma_a)}{\partial \gamma_a} + \omega_i \frac{\partial E_p(x,\gamma)}{\partial \gamma_a} - \mu_a \lambda_a; a \in L^i; i = 1, \dots, I$
$F^3(X) = \mu_a \gamma_a - \sum_{q \in P} x_q \delta_{aq}; a \in L^i; i = 1, \dots, I$

Figure 4: Sustainable Freightage Network VI

The corresponding variational inequality model presented in Figure 4 is defined in terms of the product flows (x) on each unique path p from firm i to demand market k as well as the operation frequencies (γ) and Lagrange multipliers (λ) associated with each link a. Total operational costs, frequency of operation costs, emission costs, demand functions, and link-path indicator functions are designated by C, g, E, ρ, and δ_{ap} respectively.

Figures 5 and 6 reveal the performance gains achieved in employing Runge-Kutta methods, specifically Heun-Euler with $\Delta_0 = 10^{-1}$ (RKHE) and Cash-Karp with $\Delta_0 = 10^{-3}$ (RKCK), over Euler's method and the well-known extragradient (EG) method in determining the solution to the above VI as the size of the network (dimensionality of X) grows. Since in general, the value of our gap function in this example, $g_\alpha(X)$, grows with network size, we elect to judge convergence by the reduction in $g_\alpha(X)$ from the first iteration, $g_\alpha(X)/g_\alpha(X_0) < \epsilon, \epsilon = 10^{-6}$. As shown in the figures, the

Runge-Kutta methods scale better than the other two methods both in terms of number of iterations to convergence and time to completion. Even with constant iterations, runtime increases primarily because the evaluation time of the mapping $F(X)$, F_{eval}, increases by more than 100 times over the growth of the networks.

Figure 5: This figure compares our proposed adaptive stepsize RK methods against Euler's method (Algorithm 1 where $D = I$) used in (Nagurney, Yu, and Floden 2013) and the well-known extragradient method on the basis of iteration count.

Figure 6: This figure repeats the comparison in Figure 5 on the basis of runtime.

5.2 "Next Generation" Internet Experiment

In this second experiment, providers compete to maximize profits by adjusting the quantity, quality, and price of ser-

vices delivered (Nagurney and Wolf 2014). The game described in Section 2.2 is modeled with the variational inequality in Figure 7. Like the supply chain model, we assume the governing equilibrium is Cournot-Nash-Bertrand and the utility functions are all concave and fully differentiable.

$$\langle F(X^*), X - X^* \rangle \geq 0, \forall X \in \mathcal{K}, \text{ where } X = (Q, q, \pi) \in \mathbb{R}_+^{3mno}$$
$$F_{ijk}^1(X) = \frac{\partial f_i(Q)}{\partial Q_{ijk}} + \pi_{ijk} - \rho_{ijk} - \sum_{h=1}^n \sum_{l=1}^o \frac{\partial \rho_{ihl}(Q,q)}{\partial Q_{ijk}} \times Q_{ihl}$$
$$F_{ijk}^2(X) = \sum_{h=1}^m \sum_{l=1}^o \frac{\partial c_{hjl}(Q,q)}{\partial q_{ijk}}$$
$$F_{ijk}^3(X) = -Q_{ijk} + \frac{\partial oc_{ijk}(\pi_{ijk})}{\partial \pi_{ijk}}$$

Figure 7: Service-Oriented Internet

The variational inequality in Figure 7 is defined in terms of the service quantity (Q), quality (q), and price (π) delivered from service provider i by network provider j to consumer k. Production costs, demand functions, delivery costs, and delivery opportunity costs are designated by f, ρ, c, and oc respectively.


Figure 8: This figure compares our proposed adaptive stepsize RK methods against Euler's method (Algorithm 1 where $D = I$) used in (Nagurney and Wolf 2014) and the well-known extragradient method on the basis of iteration count.

Figures 8 and 9 repeat the same ϵ-convergence ($\epsilon = 10^{-6}$) experiment for the internet model as Figures 5 and 6 did for the freightage network. Here, RKHE is run with $\Delta_0 = 10^{-2}$ and RKCK with $\Delta_0 = 10^{-5}$. As shown in the figure, RKHE and RKCK scale better than the other two methods both in terms of number of iterations to convergence and time to completion. Here, F_{eval} increases by over 200 times.

6 Conclusion

In this paper, we explored a computational game theory framework in AI based on variational inequalities. We analyzed two real-world domains, both involving oligopolistic
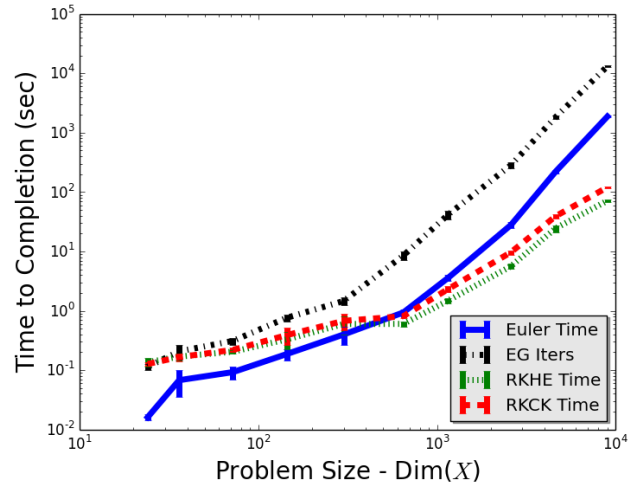

Figure 9: This figure repeats the comparison in Figure 8 on the basis of runtime.

network economies. We proposed a novel Runge Kutta algorithmic framework to solve such networked VIs, and showed that it scales far better than standard popular algorithms for solving VIs, such as the projection method and the extragradient method. The solutions to these variational inequality formulations contain rich information useful for improving the steady-states of the corresponding networks. For example, emissions regulations would see much less resistance if we could convince businesses that improving their shipping fleets and cutting emissions could result in actual increases in profits. Similarly, equilibrium solutions to the service-oriented internet would tell us how we could shape future internet policies to ensure a more comprehensive economic infrastructure.

References

Dafermos, S. 1980. Traffic equilibria and variational inequalities. *Transportation Science* 14:42–54.

Dutta, J. 2012. When is a gap function good for error bounds? Technical report, Optimization Online.

Facchinei, F., and J., P. 2003. *Finite-Dimensional Variational Inequalities and Complimentarity Problems*. Springer.

Fudenberg, D., and Levine, D. K. 1998. The theory of learning in games.

Hartman, P., and Stampacchia, G. 1966. On some nonlinear elliptic differential functional equations. *Acta Mathematica* 115:271–310.

Juditsky, A.; Nemirovski, A.; et al. 2011a. First order methods for nonsmooth convex large-scale optimization, i: general purpose methods. *Optimization for Machine Learning* 121–148.

Juditsky, A.; Nemirovski, A.; et al. 2011b. First order methods for nonsmooth convex large-scale optimization, ii: uti-

lizing problems structure. *Optimization for Machine Learning* 149–183.

Korpelevich, G. 1977. The extragradient method for finding saddle points and other problems. *Matekon* 13:35–49.

Liu, B.; Mahadevan, S.; and Liu, J. 2012. Regularized off-policy td-learning. In *Advances in Neural Information Processing Systems*, 836–844.

Nagurney, A., and Wolf, T. 2014. A cournot–nash–bertrand game theory model of a service-oriented internet with price and quality competition among network transport providers. *Computational Management Science* 11(4):475–502.

Nagurney, A., and Zhang, D. 1996. *Projected Dynamical Systems and Variational Inequalities with Applications*. Kluwer Academic Press.

Nagurney, A.; Yu, M.; and Floden, J. 2013. Supply chain network sustainability under competition and frequencies of activities from production to distribution. *Computational Management Science* 10(4):397–422.

Nagurney, A. 1998. *Network economics: A variational inequality approach*, volume 10. Springer.

Nagurney, A. 1999. *Network Economics: A Variational Inequality Approach*. Kluwer Academic Press.

Nesterov, Y. 2007. Dual extrapolation and its application to solving variational inequalities and related problems. *Mathematical Programming Series B.* 109:319–344.

Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*. Cambridge University Press.

Peng, J., and Yao, J. 2008. A new hybrid extragradient method for generalized mixed equilibrium problems, fixed point problems, and variational inequality problems. *Taiwanese Journal of Mathematics* 12:1401–1432.

Sen, A., and Shanno, D. F. 2008. Optimization and dynamical systems algorithms for finding equilibria of stochastic games. *Optimisation Methods & Software* 23(6):975–993.