Control, Intervention, and Behavioral Economics over Human Social Networks against COVID-19

M. Nagahara[a] , B. Krishnamachari[b], M. Ogura[c], A. Ortega[b], Y. Tanaka[d], Y. Ushifusa[e], and T. W. Valente[f]

[a]Institute of Environmental Science and Technology, The University of Kitakyushu, Japan; [b]Department of Electrical and Computer Engineering, University of Southern California, USA; [c]Graduate School of Information Science and Technology, Osaka University, Japan; [d]Department of Electrical Engineering and Computer Science, Tokyo University of Agriculture and Technology, Japan; [e]Faculty of Economics and Business Administration, The University of Kitakyushu, Japan; [f]Department of Preventive Medicine, University of Southern California, USA;

ARTICLE HISTORY
Compiled July 16, 2021

ABSTRACT
In this short paper, we propose a new direction of cross-cutting research for prediction and control of spreading COVID-19 viruses over a human social network. Such a network consists of human agents whose behaviors are highly uncertain and biased. To predict and control such an uncertain network, we need to employ various researches such as control theory, signal processing, machine learning, and behavioral economics. In this article, we introduce our recent research results, and propose future research topics to overcome the COVID-19 pandemic.

KEYWORDS
optimal control, human social networks, COVID-19, cyber-physical-human system (CPHS), behavioral economics, causal inference, multi-armed bandits

1. Introduction

Optimal control is a fundamental theory in robotics. It gives mathematically strict ways of designing control systems, for example, a robotic manipulator with minimum time to achieve a given task [1], stabilizing bilateral teleoperators under network time delays [2], and designing human-like dynamic running motion [3]. The performance achievable by optimal control highly depends on the accuracy of the mathematical model used to characterize the object to be controlled. Thus, optimal control has been mainly applied to physical systems such as mechanical and electrical systems that can be modeled very precisely.

Optimal control is also being applied to *human social networks*, viewed as *cyber-physical-human systems* (CPHS) [4], and in particular such ideas have been discussed in the context of the *COVID-19* pandemic [5]. Controlling these networks is not easy since they consist of human agents whose behaviors are highly uncertain. Thus, we cannot directly apply conventional optimal control to these systems by treating people

CONTACT M. Nagahara. Email: nagahara@ieee.org

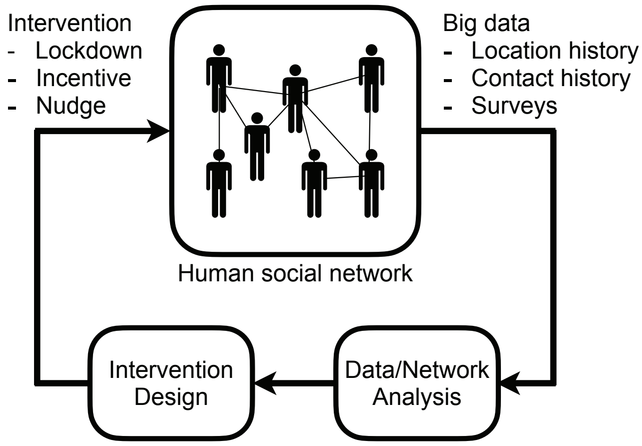

Figure 1. Cyber-physical-human system for control of virus spreading

in the network as physical systems such as electric motors. Instead, we need to design effective *interventions* to change human behaviors and achieve a global goal within the CPHS, such as reducing the infection rate in the network. This is not only a control problem but also an important problem in *behavioral economics*. In this short paper, we propose a new direction of cross-cutting research: *control, intervention, and behavioral economics over networks*.

Figure 1 illustrates the cyber-physical-human system we consider in this paper. The "Data/Network Analysis" block in Figure 1 constructs network models from observed data. This leverages the fact that it now has become possible to collect massive amounts of data about human behaviors from various sources, such as open online surveys [7] and sensor measurements. Methods to estimate networks from these data have been widely studied in the fields of signal processing and machine learning [8–14]. These methods rely on mathematical priors for the data on underlying network structures (e.g., signal smoothness on the network), and then, they estimate the most likely (potentially time-varying) network from observed data.

Next, based on network models learned from data, we design interventions to "control" the network ("Intervention Design" block in Figure 1). The main purpose of this paper is to propose an intervention design, for which we adopt three approaches:

- *model-based* optimal control for social distancing,
- *experiment-based* incentive design with causal inference,
- and *learning-based* incentive design with multi-armed bandits (MAB).

These methods are related to each other. For example, the model-based approach can give a theoretical limit of achievable performance of control to the experiment- and learning-based approaches.

Model-based approach: By constructing networked epidemic models from observed data, we can predict the propagation of viruses in human social networks for designing effective control strategies. To achieve this goal, we can leverage the results from the emerging research field of *network epidemiology* [15], which provides various frameworks and tools for modeling, analysis, and control of epidemic processes taking place in complex networks. One of the trends in this field is to utilize tools from systems and control theory for better understanding epidemic spreading processes [16]. For example, we have recently shown that the time-variability of human contact networks can be efficiently captured [17] by using the theory of Markov jump linear systems, a model of stochastic dynamical systems studied in the systems and control theory [18]. We have also shown that modeling of epidemic processes as positive systems [19] allows us to use geometric programming [20] to design effective interventions on epidemic processes [21].

Experience-based design: In behavioral economics, *causal inference* is a powerful tool for incentive design [22] for estimation of effectiveness of incentives. More recently, the method has been applied to engineering problems. For example, a recent study [6] reports an application of behavioral economics to energy management systems. In this research, we investigated the effectiveness of monetary and non-monetary incentives through a field experiment.

Learning-based design: Apart from the two approaches mentioned above, we can also adopt analysis and design by simulation and learning. In one ongoing work, we are simulating the epidemic spread on a large college campus, using data about student classroom enrollments to determine which individuals (students and teachers) are likely to encounter each other, when, and where. Taking into account a model for indoor airborne dispersion of virus particles and consequent infection, we are able to analyze how the number of cases would grow over the course of a week, and how the growth of infections is affected by factors such as mask-wearing, and reduced occupancy of courses [23]. Another opportunity that presents itself in this context is to leverage machine learning algorithms, such as multi-armed bandits to design incentives that are effective at reducing contacts and spreading of infections.

Finally, we mention some recent researches on social networks and health. Social networks are significant influences on a wide range of behaviors. Contact networks provide the vector for disease spread [24], and throughout history have changed the course of human development [25,26]. Social networks have also been demonstrated to have strong influences on many health behaviors [27] such as obesity [28], adolescent tobacco use [29], contraceptive choices [30], physician behavior [31], and country health policy adoption [32]. The pathways and mechanisms for these influences vary and include imitation, persuasion, peer pressure, modeling, and so on. In general, the associations between individual health behaviors and that of one's peers are quite strong, rivaling the effects of socio-economic status.

2. Intervention over networks

To control the spread of viruses over a human social network, we need to design interventions that can be effective in changing human behaviors in the network. A lockdown is an extreme example of a possible intervention that drastically reduces connections, which leads to the reduction of infection rate, but at the same time, negatively (and significantly) impacts the local and global economy.

Decision makers have many interventions to choose from, and these have a wide variety of strictness, frequency and incentives. Beyond simulations, interventions are required to be effective in the real world. As mentioned, human behaviors are highly uncertain and biased, and thus developing interventions (control strategies) that work well, requires using both engineering and social science tools and concepts.

Therefore, we propose novel design and analysis methods of interventions/incentives, based on model-based (Section 2.1), experiment-based (Section 2.2), and learning-based (Section 2.3) methods as mentioned in the previous section.

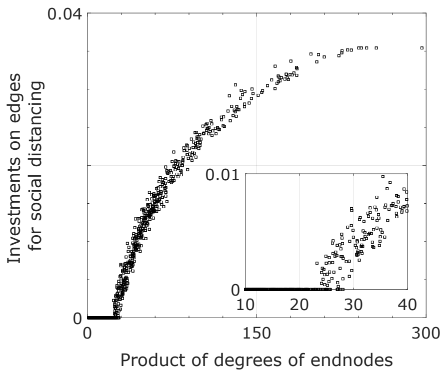

Figure 2. A threshold phenomenon in the optimal investments on edges for social distancing. Horizontal axis: the centrality of an edge measured by the product of degrees of its endnodes. Vertical axis: the optimal investment on an edge for promoting social distancing. The optimal investment found via convex optimization [21] are zero for edges with centrality below about 20 (see the inset figure). The underlying network structure is created by the Barabási-Albert model.

2.1. Optimal social distancing

Although vaccines are coming out and governments are working hard to distribute them as fast as possible, the main tool being used to contain and suppress the current pandemic is still mostly social distancing [33,34], in which healthy individuals avoid contact with infected (and potentially infected) individuals in order to protect themselves against the disease spread. Social distancing would also be a key major tool for coping with future pandemics of emerging diseases, as it is not realistic to prepare vaccines for unknown diseases in advance. However, the current social distancing strategies are often performed without a solid scientific evidence, and the quantitative evaluation of the strategies remains a difficult task. For this reason, a solid mathematical approach to design an effective social distancing strategy is required to develop effective responses to the current pandemic and to improve preparedness for similar situations in future. In this context, we have been working to identify how a limited resource should be allocated over a social network to encourage social distancing for the optimal containment of epidemics [21,35]. An interesting finding is the existence of a threshold phenomenon, where the optimal resource investments for social distancing are zero for edges whose edge-centrality is below a certain number (see Figure 2 for an illustration). However, these preliminary results focus on the most basic epidemic model called the susceptible-infected-susceptible (SIS) model, and assume a complete knowledge of the network structure.

Continuing our research effort, we are planning to generalize these preliminary results to more realistic epidemic models such as the SIDARTHE model [36] by building on our attempt to analyze the susceptible-infected-recovered (SIR) model [37]. In order to account for the unavoidable uncertainty in modeling human social networks, we will further incorporate our ongoing research on robust control for positive systems [38]. The developed mathematical tools will be applied to, for example, real time feedback control of crowd density in public areas. To realize such control, we will need to learn from the information available about current crowd density to develop a model of the disease spread so that optimal social distancing can be achieved by interventions such as temporarily restricting entrance to certain areas. We are also planning to develop a mathematical framework to limit the frequency of interventions by leveraging our ongoing research on event-triggered control [39] and sparse control [40–42].

Yet another important direction of research is to incorporate the new-normal social distancing strategies that societies have taken in the current pandemic [43,44] into our theoretical development. The diversity in the strategies would urge us to develop flexible frameworks able to accommodate the various types of temporal changes in the structure of social networks.

2.2. Causal inference and field experiments based on behavioral economics

As mentioned above, a cyber-physical-human system (CPHS) consists of humans as agents whose behaviors are observed from sensors through the Internet. An example of CPHS is an urban traffic signal control system based on networked vehicle data from Internet-connected sensors [45]. A challenging problem is to estimate and control such a CPHS with human behaviors that are highly uncertain and biased. To solve this problem, we adapt methods in behavioral economics for analyzing interventions (monetary and non-monetary incentives). A novelty of this research is to design interventions that can change the behaviors of humans who interact with each other in the network, an approach we call *network-aware design of interventions*.

Our ongoing work studies the impacts of monetary and non-monetary incentives on energy-saving behaviors using a field experiment conducted in Japan [6]. In this research, a novel method of causal inference, the *causal forest*, was proposed and the difference between monetary and non-monetary incentives was revealed.

We will extend this method to develop network-aware causal inference and estimate the effect of interventions in the context of epidemic control. In particular, we will investigate how different incentives work for various types of networks with different space and time scales. For example, we plan a field experiment in Kanmon Strait Museum in Kitakyushu by using the social distancing detection system developed by Mishima OA Systems [46]. We will use monetary incentives (e.g., reduction of entrance fees or museum gifts) and non-monetary incentives (e.g., showing a heat map on a digital signage set on the museum floor) to encourage behavioral change. We then validate the effectiveness of these incentives for changing the network state (e.g., the maximum density in the heat map, the maximum degree of the network, and sparsity).

2.3. Online learning for individualized incentives

Online learning techniques, particularly multi-armed bandits (MAB), offer a way to learn the best choices in an environment that yields stochastic feedback. In the classic MAB formulation, there are n arms, each of which yields a random reward from an unknown distribution [47]. The player must select an arm at each time based on prior observations, with the goal of maximizing the expected total reward over time. The classic metric for performance is called regret, measured in case of independent arms as the gap between the rewards collected by a player following a particular policy and that obtained by a genie that has knowledge of the arm reward distributions. In case of independent arms, it is known that the cumulative regret of the best policy grows logarithmically with time, which implies that the time-averaged regret goes to 0 asymptotically (i.e., over time, the player is able to get a time-averaged reward that is the same as the genie). Several policies with provable logarithmic regret have been designed [47–49].

Building on our prior work on designing and analyzing algorithms and novel applications of multi-armed bandits including extensions to combinatorial and network settings [50–53], we propose to formulate the problem of incentivizing people in a community in a similar way. Consider a given set of incentives. In the simplest formulation, the goal is to identify which particular incentive is the most effective at maximizing the desired benefit (e.g., with respect to ensuring that people stay home, minimizing contact with others during high epidemic risk times), by applying different incentives over time, observing their impact (a form of stochastic feedback given variations in individual responses) and adapting to spend more time using the most effective incentives while still periodically sampling the less effective incentives to account for the possibility of misleading samples in light of the stochastic feedback. A more sophisticated formulation could yield as an output a complex combinatorial collection of incentives, each personalized to different clusters of individuals; this could be posed as a joint clustering and contextual-bandit problem. We plan to explore the practical design of such an online learning-based incentive system by investigating how to incorporate feedback and observations through sensors as well as manual inputs or survey responses from users.

2.4. Stability

Since the system proposed in this paper is a feedback control system including humans in the loop, we need to analyze and guarantee the *global stability* of the system. To do this, we can adapt the *passivity-based approach* [54] to our cyber-physical-human system. It is well-known that the feedback connection of two passive systems is again passive. That is, we only need to know if the systems are both passive, even if the systems are highly uncertain and difficult to obtain the precise models.

3. Conclusion

In this short paper, we have discussed a new direction of cross-cutting research to overcome the COVID-19 pandemic. We proposed the design of interventions based on mode-based optimal control design, experiment-based causal inference for effectiveness analysis of incentives, and learning-based incentive design. To overcome the COVID-19, cross-cutting collaboration is necessary. For this, other technologies in robotics are also important. For example, social robots can be used as implementation of effective incentives that change human behaviors in the cyber-physical-human system. This idea has been already proposed in e.g., [55]. The important point to use robots in the COVID-19 epidemic is that robots are free from virus infections. This viewpoint has been discussed in recent papers [56,57].

Also, our approach considers the human-social network as an "input-output" system (or a "blackbox"). If we can observe and utilize the internal states of humans, e.g., mental states, positive/negative, aggressive/defensive, the control performance may be significantly improved, as in the state-space approach in control [58].

Acknowledgement

This work was partly supported by JSPS KAKENHI JP20H02172, JP20K21008, JP19H02301, JP18K13777, 20H02145, JST PRESTO JPMJPR1935, and JST RISTEX JPMJRX19I2.

References

[1] Bobrow J, Dubowsky S, Gibson J. Time-optimal control of robotic manipulators along specified paths. The International Journal of Robotics Research. 1985;4(3):3–17.

[2] Leung GMH, Francis BA, Apkarian J. Bilateral controller for teleoperators with time delay via μ-synthesis. IEEE Trans Robot Autom. 1995;11(1):105–116.

[3] Schultz G, Mombaur K. Modeling and optimal control of human-like running. IEEE/ASME Trans Mechatronics. 2010;15(5):783–792.

[4] Sowe SK, Simmon E, Zettsu K, et al. Cyber-physical-human systems: Putting people in the loop. IT Professional. 2016;18(1):10–13.

[5] Annaswamy A. Cyberphysical and human systems: A new frontier in control systems [president's message]. IEEE Control Syst Mag. 2020;40(4):8–9.

[6] Murakami K, Shimada H, Ushifusa Y, et al. Heterogeneous treatment effects of nudge and rebate: Causal machine learning in a field experiment on electricity conservation. Graduate School of Economics Discussion Paper Series, Kyoto University. 2020;issue.E-20-003:1–46.

[7] Ojo O, García-Agundez A, Girault B, et al. Coronasurveys: Using surveys with indirect reporting to estimate the incidence and evolution of epidemics. arXiv preprint arXiv:200512783. 2020;.

[8] Dong X, Thanou D, Rabbat M, et al. Learning graphs from data: A signal representation perspective. IEEE Signal Processing Magazine. 2019;36(3):44–63.

[9] Mateos G, Segarra S, Marques AG, et al. Connecting the dots: Identifying network structure via graph signal processing. IEEE Signal Processing Magazine. 2019;36(3):16–43.

[10] Kalofolias V. How to learn a graph from smooth signals. In: th International Conference on Artificial Intelligence and Statistics AISTATS, Cadiz, Spain; 2016.

[11] Hallac D, Park Y, Boyd S, et al. Network Inference via the Time-Varying Graphical Lasso. In: Proc. of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17), New York, NY, USA; 2017. pp. 205–213.

[12] Egilmez HE, Pavez E, Ortega A. Graph learning from data under laplacian and structural constraints. IEEE Journal of Selected Topics in Signal Processing. 2017;11(6):825–841.

[13] Yamada K, Tanaka Y, Ortega A. Time-varying graph learning based on sparseness of temporal variation. In: Proc. Int. Conf. Acoust., Speech, Signal Process. (ICASSP); 2019. pp. 5411–5415.

[14] Yamada K, Tanaka Y, Ortega A. Time-varying graph learning with constraints on graph temporal variation. arXiv preprint arXiv:200103346. 2020;.

[15] Pastor-Satorras R, Castellano C, Van Mieghem P, et al. Epidemic processes in complex networks. Reviews of Modern Physics. 2015;87(3):925–979.

[16] Nowzari C, Preciado VM, Pappas GJ. Analysis and control of epidemics: A survey of spreading processes on complex networks. IEEE Control Systems. 2016;36(1):26–46.

[17] Ogura M, Preciado VM. Stability of spreading processes over time-varying large-scale networks. IEEE Transactions on Network Science and Engineering. 2016 1;3(1):44–57.

[18] Costa OLV, Fragoso MD, Todorov MG. Continuous-time Markov Jump Linear Systems. Springer; 2013.

[19] Farina L, Rinaldi S. Positive Linear Systems: Theory and Applications. Wiley-Interscience; 2000.

[20] Boyd S, Kim SJ, Vandenberghe L, et al. A tutorial on geometric programming. Optimization and Engineering. 2007;8(1):67–127.

[21] Ogura M, Preciado VM. Epidemic processes over adaptive state-dependent networks. Physical Review E. 2016;93:062316.

[22] G. W. Imbens and D. B. Rubin. Causal inference for statistics, social, and biomedical science. Cambridge University Press; 2015.

[23] Hekmati A, Luhar M, Krishnamachari B, et al. Simulation based analysis of covid-19 spread through classroom transmission on a university campus. in submission. 2021;.

[24] Anderson RM, May RM. Infectious diseases of humans: Dynamics and control. New York, NY: Oxford University Press; 1991.
[25] Diamond J. Guns, germs and steel: The fates of human societies. New York: Norton; 1997.
[26] McNeill WH. Plagues and peoples. New York: Doubleday; 1977.
[27] Valente TW. Social networks and health: Models, methods, and applications. New York: Oxford University Press; 2010.
[28] Christakis NA, Fowler JH. The spread of obesity in a large social network over 32 years. The New England Journal of Medicine. 2007;357:370–379.
[29] Valente T, Fujimoto K, Soto D, et al. A comparison of peer influence measures as predictors of smoking among predominately hispanic/latino high school adolescents. Journal of Adolescent Health. 2012;52:358–364.
[30] Valente TW, Watkins S, Jato MN, et al. Social network associations with contraceptive use among cameroonian women in voluntary associations. Social Science and Medicine. 1997;45:677–687.
[31] Iyengar R, Van den Bulte C, Valente TW. Opinion leadership and contagion in new product diffusion. Marketing Science. 2011;30:195–212.
[32] Valente TW, Dyal SP, Chu KC, et al. Diffusion of innovations theory applied to global tobacco control treaty ratification. Social Science & Medicine. 2015;145:89–97.
[33] Bell D, Nicoll A, Fukuda K, et al. Nonpharmaceutical interventions for pandemic influenza, national and community measures. Emerging Infectious Diseases. 2006;12(1):88–94.
[34] Funk S, Salathé M, Jansen VAA. Modelling the influence of human behaviour on the spread of infectious diseases: a review. Journal of The Royal Society Interface. 2010 9; 7(50):1247–1256.
[35] Ogura M, Preciado VM, Masuda N. Optimal containment of epidemics over temporal activity-driven networks. SIAM Journal on Applied Mathematics. 2019;79(3):986–1006.
[36] Giordano G, Blanchini F, Bruno R, et al. Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. Nature Medicine. 2020 6;26(6):855–860.
[37] Ogura M, Preciado VM. Efficient containment of exact SIR Markovian processes on networks. In: 55th IEEE Conference on Decision and Control; 2016. p. 967–972.
[38] Ogura M, Kishida M, Lam J. Geometric programming for optimal positive linear systems. IEEE Transactions on Automatic Control. 2020 11;65(11):4648–4663.
[39] Onoue Y, Hashimoto K, Ogura M, et al. Event-triggered control for mitigating SIS spreading processes. {\upshape Submitted for publication}. 2020;.
[40] Nagahara M, Quevedo DE, Nešić D. Maximum hands-off control: a paradigm of control effort minimization. IEEE Trans Autom Control. 2016;61(3):735–747.
[41] Nagahara M, Chatterjee D, Challapalli N, et al. CLOT norm minimization for continuous hands-off control. Automatica. 2020;113:108679.
[42] Nagahara M. Sparsity methods for systems and control. Now Publishing; 2020.
[43] Flaxman S, Mishra S, Gandy A, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. Nature. 2020;584(7820):257–261.
[44] Sebhatu A, Wennberg K, Arora-Jonsson S, et al. Explaining the homogeneous diffusion of COVID-19 nonpharmaceutical interventions across heterogeneous countries. Proceedings of the National Academy of Sciences of the United States of America. 2020; 117(35):21201–21208.
[45] Su R. Real-time V2X-based traffic signal control for large traffic networks. IEEE CDC 2019 Workshop: Systems and Control for Smart Society and Cyber-Physical & Human Systems. 2019;.
[46] Mishima OA Systems. Social distance AI ; 2020. https://www.mishimaoa.co.jp/products/sds-ai/.
[47] Lai TL, Robbins H. Asymptotically efficient adaptive allocation rules. Advances in applied mathematics. 1985;6(1):4–22.

[48] Auer P, Cesa-Bianchi N, Fischer P. Finite-time analysis of the multiarmed bandit problem. Machine learning. 2002;47(2):235–256.
[49] Kaufmann E, Korda N, Munos R. Thompson sampling: An asymptotically optimal finite-time analysis. In: International conference on algorithmic learning theory; Springer; 2012. p. 199–213.
[50] Ahmad SHA, Liu M, Javidi T, et al. Optimality of myopic sensing in multichannel opportunistic access. IEEE Transactions on Information Theory. 2009;55(9):4040–4050.
[51] Gai Y, Krishnamachari B, Jain R. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. IEEE/ACM Transactions on Networking. 2012;20(5):1466–1478.
[52] Mansourifard P, Jazizadeh F, Krishnamachari B, et al. Online learning for personalized room-level thermal control: A multi-armed bandit framework. In: Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings; 2013. p. 1–8.
[53] Sakulkar P, Krishnamachari B. Stochastic contextual bandits with known reward functions. arXiv preprint arXiv:160500176. 2016;.
[54] Hatanaka T, Chopra N, Fujita M, et al. Passivity-based control and estimation in networked robotics. Springer; 2015.
[55] Kanda T, Hirano T, Eaton D, et al. Interactive robots as social partners and peer tutors for children: A field trial. Human–Computer Interaction. 2004;19(1-2):61–84.
[56] Zeng Z, Chen PJ, Lew AA. From high-touch to high-tech: COVID-19 drives robotics adoption. Tourism Geographies. 2020;22(3):724–734.
[57] Taniguchi T, Hafi L, Hagiwara Y, et al. Semiotically Adaptive cognition: Toward the realization of remotely-operated service robots for the new normal symbiotic society. Advanced Robotics. 2021;.
[58] H. K. Khalil. Nonlinear systems. Prentice Hall; 2002.