Noncooperatively Optimized Tolerance: Decentralized Strategic Optimization in Complex Systems

Yevgeniy Vorobeychik, Jackson R. Mayo, Robert C. Armstrong, and Joseph R. Ruthruff

Phys. Rev. Lett. **107**, 108702 — Published 1 September 2011

DOI: 10.1103/PhysRevLett.107.108702

Noncooperatively Optimized Tolerance: Decentralized Strategic Optimization in Complex Systems

Yevgeniy Vorobeychik, Jackson R. Mayo, Robert C. Armstrong, and Joseph R. Ruthruff
Sandia National Laboratories, P.O. Box 969, Livermore, California 94551-0969, USA

We introduce noncooperatively optimized tolerance (NOT), a game theoretic generalization of highly optimized tolerance (HOT), which we illustrate in the forest fire framework. As the number of players increases, NOT retains features of HOT, such as robustness and self-dissimilar landscapes, but also develops features of self-organized criticality (SOC). The system retains considerable robustness even as it becomes fractured, due in part to emergent cooperation between players, and at the same time exhibits increasing resilience against changes in the environment, giving rise to intermediate regimes where the system is robust to a particular distribution of adverse events, yet not very fragile to changes.

Introduction Highly optimized tolerance (HOT) and self-organized criticality (SOC) have received considerable attention as alternative explanations of emergent power-law cascade distributions [1, 2]. The SOC model [1, 3, 4] posits that systems can naturally arrive at criticality and power-law cascades, independently of initial conditions, by following simple rule-based processes. Among the important features of SOC are (a) self-similarity and homogeneity of the landscape, (b) fractal structure of cascades, (c) a small power-law exponent (i.e., heavier tails), and (d) low density and low yield (e.g., in the context of the forest fire model, described below). HOT [2, 5–7], in contrast, models complex systems that emerge as a result of optimization in the face of persistent threats. While SOC is motivated by largely mechanical processes, the motivation for HOT comes from evolutionary processes and deliberately engineered systems, such as the electric power grid. The key features of HOT are (a) a highly structured, self-dissimilar landscape, (b) a high power-law exponent, and (c) high density and high yield [6].

HOT and SOC can be cleanly contrasted in the context of the forest fire model [2, 4], which features a grid, usually two-dimensional, with each cell being a potential site for a tree. Intermittently, lightning strikes one of the cells according to some probability distribution. If there is a tree in the cell, it is set to burn. At that point, a cascade begins: fires spread recursively from cells that are burning to neighboring cells that contain trees, engulfing the entire connected component in which they begin (in our implementation, fires wrap around the grid walls, so there are effectively no boundaries). In the classical forest fire model (SOC) a tree sprouts in every empty cell with some fixed probability p. In contrast, the HOT model conceives of a global optimizer choosing the configuration of each cell (i.e., whether a tree will grow or not); what emerges globally as a consequence is a collection of large connected components of trees separated by "barriers" of no trees. The HOT model is deliberately robust to lightning strikes with the specified distribution; however, it is also extremely fragile to changes in the lightning distribution, whereas SOC does not exhibit such fragility. The HOT landscape tends to have a highly non-uniform distribution of "fire breaks", or areas where no trees are planted, whereas the SOC landscape is homogeneous.

A natural criticism of the HOT paradigm is that, in complex systems, it is difficult to conceive of a single designer that manages to optimally design such a system. As a partial response, much work demonstrates that HOT yields qualitatively similar results when heuristic optimization or an evolutionary process is used [5, 8]. Still, most complex systems are not merely difficult to design globally, but are actually *decentralized*, with many entities responsible for parts of the whole system. Each entity is generally not motivated by global concerns, responding instead to individual incentives. For example, the Internet is fundamentally a combination of autonomous entities making their own decisions about network topology, protocols, and composition.

Our central contribution is to model complex systems as complex patterns of strategic interactions among self-interested players making independent decisions. We conceive that out of *strategic interactions* of such self-interested players emerges a system that is optimized *jointly* by all players, rather than *globally* by a single "engineer". Thus, we call our model *noncooperatively optimized tolerance* (NOT). Formally, our model is game theoretic, and we seek to characterize emergent properties of the system in a Nash equilibrium.

A Game Theoretic Forest Fire Model Suppose that each player controls a portion of a complex system and is responsible for engineering his "domain of influence" against perceived threats. The interests of different players may be opposed if, say, an action that is desirable for one has a negative impact on another. Such interdependencies (commonly referred to as *externalities*) form a central aspect of our model.

We implement the game theoretic conception of complex system engineering in the familiar two-dimensional forest fire model. In the NOT forest fire model, each player is allotted a portion of the square grid over which he optimizes his yield less cost of planting trees. Let G_i be the set of grid cells under player i's direct control, let s_i be player i's strategy expressed as a vector s_i in which $s_{i,g} = 1$ if i plants a tree in grid cell g and $s_{i,g} = 0$ otherwise, and let $\Pr\{g = 1 \mid s, s_{i,g} = 1\}$ be the probability (with respect to the lightning distribution) that a tree planted in cell g survives a fire given the joint strategy (planting) choices of all players. Denote by s the vector of all players' choices. Since exactly one player controls each grid cell, we simplify notation and use $s_g = s_{i,g}$ where i is the player controlling grid cell g. Let $N_i = |G_i|$ be the number of grid cells under i's control and ρ_i be the density of trees planted by i. Let $Y_i(s) = \sum_{g \in G_i} \Pr\{g = 1 \mid s\} s_g$ be the yield for player i. Let c denote the cost of planting a tree. The utility of player i is then

$$u_i(s) = \sum_{g \in G_i} (\Pr\{g = 1 \mid s\} - c) s_{i,g} = Y_i(s) - cN_i\rho_i.$$

The result of joint decisions by all players is a grid that is partially filled by trees, with overall density $\rho(s)$ and overall yield $Y(s)$ given by a sum ranging over the entire grid G, i.e., $Y(s) = \sum_{g \in G} \Pr\{g = 1 \mid s\} s_g$. Let N be the number of cells in the entire grid. We then define *global utility (welfare)* as

$$W(s) = \sum_{i \in I} u_i(s) = Y(s) - cN\rho(s).$$

Note that when $m = 1$, $W(s)$ coincides with the lone player's utility. A part of our endeavor below is to characterize $W(s^*)$ and $\rho(s^*)$ when s^* is a Nash equilibrium, defined as a configuration of joint decisions by all players such that no individual player can gain by choosing an alternative strategy (planting configuration) s_i' *keeping the decisions of other players fixed*.

We systematically vary several model parameters. The first is the number of players m, which we vary from $m = 1$ to N, fixing the size of the grid at $N = 128 \times 128$. The former extreme corresponds precisely to the HOT setting, while in the latter the players are entirely myopic in their decision problems, each concerned with only a single cell of the grid. The entire range of player variation is $m \in \{1, 2^2, 4^2, 8^2, 16^2, 32^2, 64^2, 128^2\}$. The second parameter that we vary is the cost of planting trees: $c \in \{0, 0.25, 0.5, 0.75, 0.9\}$. Finally, we vary the scale of the lightning distribution, which is always a truncated Gaussian centered at the top left corner of the grid. We let the variance (of the Gaussian before truncation) be N/v, and vary $v \in \{0.1, 1, 10, 100\}$. For example, at $v = 0.1$ the distribution of lightning strikes is approximately uniform over the grid, while at $v = 100$ the distribution is highly concentrated in the top left corner. We divide the grid among m players by partitioning it into m identical square subgrids.

Analysis of the NOT Forest Fire Model Some intuition is provided by initially considering a mathematically tractable one-dimensional forest fire setting. There, while the density of planting approaches 1 in both the optimal and equilibrium configurations as N increases, it turns out that the equilibrium density is generally higher than optimal [9]. This agrees with our intuition on the consequence of negative externalities of decentralized planting decisions: when a player decides whether to plant a tree, he takes into account only the concomitant chance of his own tree burning down, and not the global impact the decision has on the sizes of cascades.

A full analysis of the two-dimensional model in all the relevant parameters is beyond mathematical tractability. Furthermore, the problem of computing exact equilibria, or even exact *optima* for any player, is intractable, as the size of the space of joint player strategies in our setting is 2^{16384}. Nevertheless, it turns out that simple iterative algorithms for approximating equilibria as well as optimal decisions by individual players are extremely effective. Specifically, we use a variant of *best response dynamics* for approximating Nash equilibria, which iteratively optimizes each player's strategy, keeping strategies of other players fixed [10]. (We found that both asynchronous and partially synchronous versions of best response dynamics yield similar results [9]; below we report on the asynchronous implementation.) Within this procedure, we approximate optimal responses of individual players using *sampled fictitious play* [11]. In sampled fictitious play, each grid cell controlled by player i becomes a "player" in a co-operative subgame (where each cell has i's utility as its goal), and random subsets of cells are iteratively chosen to make simultaneous optimizing decisions.

Our first question concerns the variation of global utility $W(s^*)$ with the number of players m, the cost c, and the parameter v governing variance of the lightning distribution. First, note that $W(s^*)$ will be no better than optimal for $m > 1$, and it seems intuitive that it is a non-increasing function of m. Additionally, when $c = 0$ and $m = N$, we have a global utility of 0, since the only equilibria involve either all, or all but one, players planting trees [9]. In the more general cases, the following simulation results are obtained [9]. When $c = 0$, the initial drop in global utility is quite shallow for $m < 256$, particularly when the lightning distribution is relatively diffuse ($v < 100$). However, once the number of players is relatively large, global utility drops dramatically, and nearly reaches 0 already when $m = 4096$. For $c > 0$, the dropoff in global utility with the number of players becomes less dramatic.

Our next task is to consider how the density changes with our parameters of interest. Based on the observation above, we expect the density to be 1, or nearly so, when $c = 0$ and $m = N$. The density should be appreciably below 1 when $m = 1$. Furthermore, the density should decrease with increasing cost c. In general, our intuition, based on all previous analysis, would suggest that density should increase with the number of players: after all, each player's decision to plant a tree does not account for the negative impact it has on other players.

Working from this intuition, the simulation results [9] are highly counterintuitive: the overall density *falls* with increasing number of players until m reaches 1024, and only when the number of players is very high (4096 and N) is it generally higher than the optimal density. This dip is especially apparent for a highly concentrated lightning distribution ($v = 100$). To understand this phenomenon we refer to Figure 1, showing actual (approximate) equilibrium grid configurations for varying numbers of players when $c = 0$ and $v = 100$. We can observe that each player's myopic self-interest induces him to construct *fire breaks* in his territory where none exist in a globally superior single-player configuration. Thus, for example, contrast Figure 1 (a) and (b). In the former, most of the grid is filled with trees, and much of the action happens in the upper left corner (the epicenter of the lightning distribution), which is filled with fire breaks that confine fires to relatively small fractions of the grid. In the latter, the upper left corner is now under the control of a single player, and other players find it beneficial to plant fire breaks of their own, since the "wasted" land amounts to only a small fraction of their land-

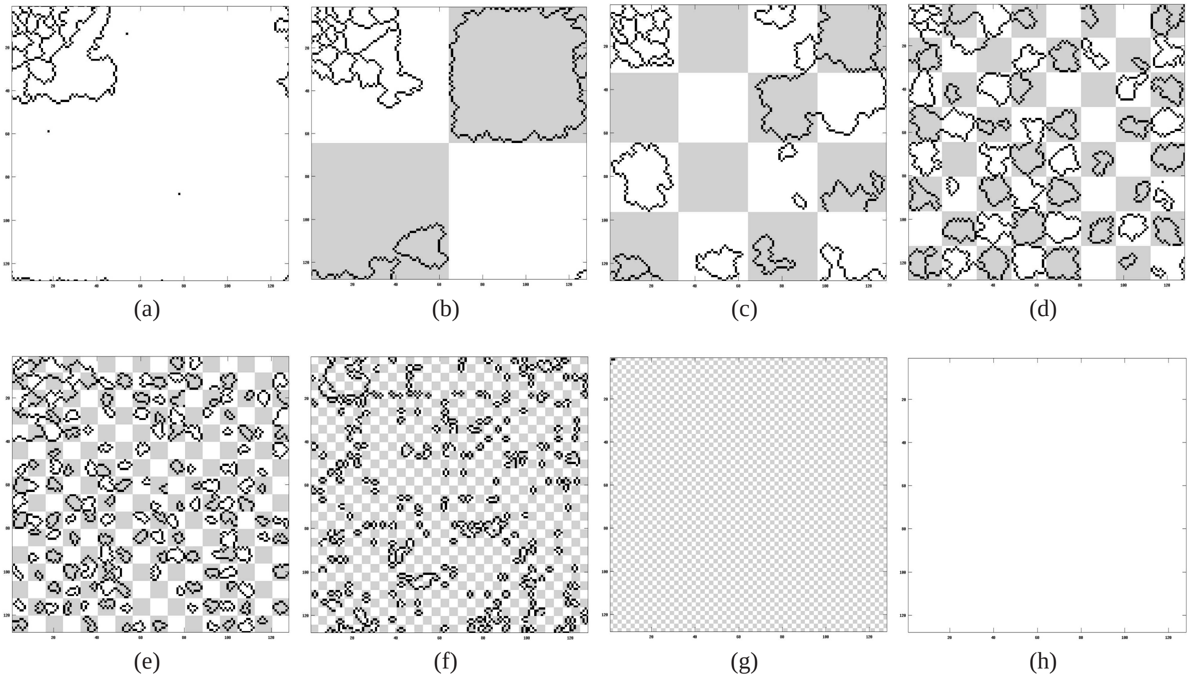

FIG. 1. Sample equilibrium grid configurations with $c = 0$, $v = 100$, and the number of players varied between 1 and $N = 16384$. Blank cells are planted and marked cells are unplanted. Player domains of influence are shaded in a checkerboard pattern. (a) 1 player, equivalent to HOT; (b) 4 players; (c) 16 players; (d) 64 players; (e) 256 players; (f) 1024 players; (g) 4096 players; (h) 16384 players. To avoid clutter, we omit the checkerboard pattern with N players, where each grid cell contains a tree. Note that players adopt different strategies in similar conditions since best response is only approximate and stochastic, and there are likely many nearly optimal configurations.

mass, and offers some protection against fire spread to the protected areas from "poorly" protected neighboring territories. With more players, we see coordination between neighbors emerge, as they jointly build mutually beneficial fire breaks, but such cooperation is not global, and becomes increasingly diffuse with greater number of players. Nevertheless, increasing the number of players results in a greater amount of total territory devoted to fire breaks by individual players or small local neighborhoods, and, as a result, an overall loss in planting density as observed.

Since the density is decreasing for intermediate numbers of players, a natural hypothesis is that the fire breaks are distributed suboptimally. We can observe this visually in Figure 1. Specifically, analysis of the equilibrium grid configurations shows that the location of fire breaks becomes less related to the lightning distribution as the number of players grows [9]. Interestingly, even for a moderate number of players ($m = 16$), the distribution of fire breaks is nearly homogeneous and almost unrelated to the lightning distribution. This suggests that global utility would remain relatively robust to changes in the lightning distribution compared to the HOT model. To verify this, we show in Figure 2 average global utility of equilibrium configuration *after the lightning distribution is randomly changed*. Whether the cost of planting trees is high or low, the figure shows significantly reduced fragility for an intermediate number of players (between 16 and 1024). Indeed, when cost is high, the system remains less fragile than HOT even in the limiting case of $m = N$. Because global utility remains relatively close to optimal across a wide range of settings when m is below 256 [9], our results suggest that the regime of intermediate numbers of players retains the robustness of HOT, while developing some features of SOC that make it less fragile to changes in the environment. Perhaps the most important reason for this phenomenon is the impact that negative externalities have on behavior of agents most susceptible to them: players closest to the epicenter of the lightning distribution tend to overplant, and others respond by building firebreaks around parts of their territory, partially protecting themselves from negative effects of neighbors' decisions. A direct consequence of these decisions is that the overall configuration remains quite robust to lightning strikes. A surprising consequence is that the resulting fire breaks form effective barriers preventing excessive spread of fire if the lightning distribution changes. When the number of players (m) is very small, however, player decisions correspond very closely to the actual lightning distribution, increasing fragility, while a very large m fragments decisions too much, and player decisions are highly myopic, with resulting configurations often not robust and highly fragile.

One of the central results of both SOC and HOT models is a power-law distribution of burnout cascades. Since our model generalizes HOT, we should certainly expect to find an approximately power-law distribution in the corresponding special case of $m = 1$. We now study how the burnout distribution behaves with respect to the parameters of interest.

Figure 3 shows fire cascade distributions on the usual log-

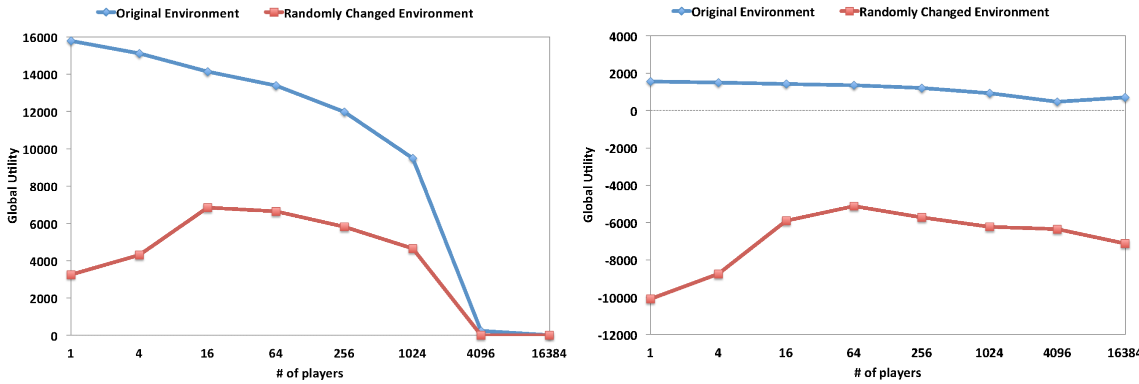

FIG. 2. Fragility of NOT configurations for $v = 100$. Given the (approximate) equilibrium configurations generated for a lightning distribution centered at the upper left corner of the grid, we changed the lightning distribution by generating the center of the Gaussian uniformly randomly from all grid locations. We then evaluated expected global utility given the altered lightning distribution. The graph plots averages of repeating this process 30-80 times, as compared to global utility for the original environment. Left: $c = 0$. Right: $c = 0.9$.

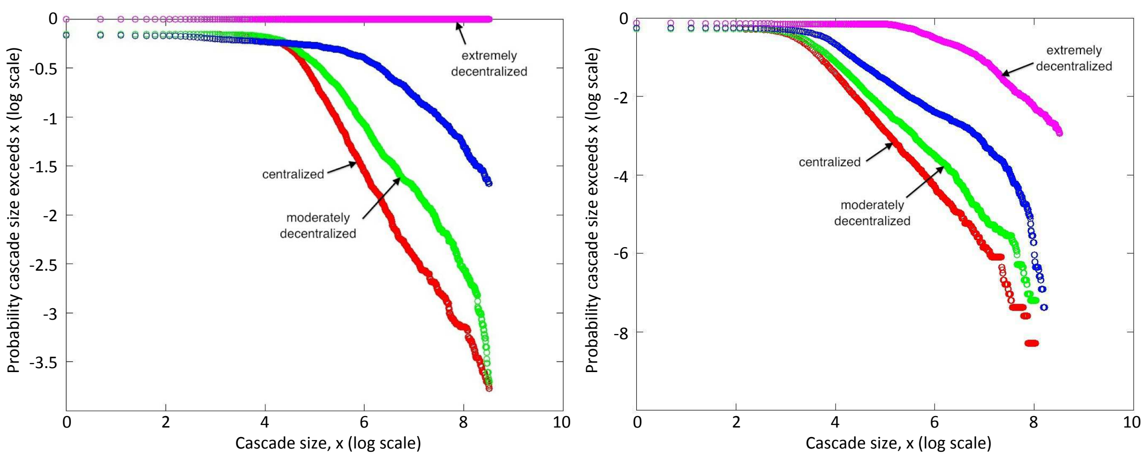

FIG. 3. Distribution of tree burnout cascades, shown on a log-log plot with $\Pr\{X \geq x\}$ on the vertical axis and x on the horizontal axis, where X is the random variable representing cascade size. The plots feature (bottom to top) $m = 1$ (red), $m = 16$ (green), $m = 256$ (blue), and $m = 4096$ (purple), with the left plot corresponding to $c = 0$ and the right plot corresponding to $c = 0.9$. Both plots correspond to $v = 10$.

log plot for $v = 10$. When $m = 1$ (red points), the results suggest an approximate power-law distribution across a range of scales. Additionally, even when m is greater than 1 but relatively small (green points), the distribution remains approximately linear across a range of scales, suggesting that the power law is likely not unique to the HOT setting. Once the number of players is large, however, the distribution of cascades less resembles a power law, and begins to feature considerable curvature even at the intermediate scales. In that sense, the NOT setting with many players is unlike both HOT and SOC. The most important aspect of the cascade distributions is that the tails are systematically increasing with the number of players in all observed settings [9].

Sandia is a multiprogram laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy under contract DE-AC04-94AL85000.

[1] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. Lett., **59**, 381 (1987).
[2] J. M. Carlson and J. Doyle, Phys. Rev. E, **60**, 1412 (1999).
[3] S. Clar, B. Drossel, and F. Schwabl, J. Phys. Cond. Matt., **8**, 6803 (1996).
[4] C. L. Henley, Phys. Rev. Lett., **71**, 2741 (1993).
[5] J. M. Carlson and J. Doyle, Phys. Rev. Lett., **84**, 2529 (2000).
[6] J. M. Carlson and J. Doyle, Proc. Natl. Acad. Sci., **99**, 2538 (2002).
[7] M. E. J. Newman, M. Girvan, and J. D. Farmer, Phys. Rev. Lett., **89**, 028301 (2002).
[8] T. Zhou, J. M. Carlson, and J. Doyle, J. Theor. Biol., **236**, 438 (2005).
[9] See EPAPS Document No. X for supporting details.
[10] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games* (The MIT Press, 1998).
[11] T.J. Lambert III, M. Epelman, and R. L. Smith, Oper. Res., **53**, 477 (2005).