2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

Games of Network Disruption and Idempotent Algorithms

William M. McEneaney Antoine Desir

Abstract— **We consider a game on the space of network disruptions. An application in command and control is used as a guide for the development of the model. The outcome of any set of physical actions depends on the information available to the controller. We suppose that information flows along a network of humans and machines. The opposing player may act to intermittently block information flow along the network. This may be combined with physical actions such that our controller might be making decisions based on information that is not as current as would be possible without the induced delays. We find that the model of information delay dynamics is best captured as a controlled min-plus linear system. We also find that the minimax value function may be represented as a min-plus convex functional over the space of delay vectors. This is a max-min linear space. Backward dynamic programming propagation of the value function leads to (max-min) sum and product compositions. This yields a particularly nice solution algorithm. Computational and solution-representation complexity are examined.**

I. Introduction

We consider a game on the space of network disruptions. An application in command and control is used as a guide for the development of the model. However, it should be clear that the problem class addressed by the theory is much larger. There is a physical domain via which the ultimate rewards are obtained. The outcome of any set of physical actions depends on the information available to the controller. We suppose that information flows along a network of humans and machines, where information processing may occur may occur at some of the nodes. The opposing player may act to intermittently block information flow along the network. This may be combined with physical actions such that our controller might be making decisions based on information that is not as current as would be possible without the disruption.

We will find that the model of information delay is best captured as a min-plus linear system. Specifically, the delay of information from node g at node γ is the minimum over the delays along available information-routes from node g to node γ. The opposing control may act to increase delays while our controller may seek to reduce such delays, leading to a game with controlled min-plus linear dynamics.

We will find that the value function may take the form of a min-plus convex functional over the space of delay vectors. Note that in this component of the analysis, we work in min-plus vector spaces. Recall that the space of standard-sense convex functions is a max-plus linear space (c.f., [10]), where finite-complexity convex functions are those given as pointwise maxima of a finite set of standard-sense affine functions. That is, they may be represented as max-plus linear combinations of standard-sense linear functions. In the work here, the standard-sense algebra is replaced by the min-plus algebra and the max-plus algebra is replaced by the max-min algebra. Consequently, min-plus convex functions are represented as max-min linear combinations of min-plus linear functions.

We will want to consider larger scale networks. Thus, the complexity of the solution representation is a critical question. We are led to questions of the complexity of max-min sums and products of finite-complexity min-plus functions. We find that this complexity grows more slowly that one would intuitively expect. We also consider max-min linear projections as a means of complexity-growth attenuation.

II. Motivational Application

We refer to figure 1 for development. Suppose there exists a network consisting of sensing nodes, action nodes, communication nodes and processing/decision nodes. An example is given in the figure. The network may be considered to be belonging to the Blue player, although mathematically this is somewhat irrelevant. The opponent is designated as the Red player. Information will flow from the Blue sensing nodes along the network. Processing and/or decisions may be made by the relevant nodes, and the results will flow to the Blue action nodes. (Of course, there would be a network associated with Red as well as with Blue, but that is outside the scope of our analysis here.) There are two levels to the game. At the physical level, the Blue action nodes (represented as groups of blue triangles in the figure), might interact with the corresponding Red action nodes. (In the figure, the Red action nodes are placed in proximity to Blue nodes in order to intuitively indicate some physical proximity, although this is unrelated to the network graph structure.) The outcome of any such interaction will depend not only on the physical state, but also on the information available to Blue at the time of such interaction. We will associate a payoff with the outcome. The dependence of outcomes on information in a purely game model is discussed in [8], and we will review that here. However, in this study, we will be concerned with a game of network disruption. Any network disruptions will induce delays in the information available to the Blue action nodes, which will degrade the value of the information. Although the game of network disruption will be the focus, some amount of discussion of the physical level game is needed in order to motivate the

Research partially supported by AFOSR.

Dept. of Mechanical and Aerospace Eng., University of California San Diego, San Diego, CA 92093-0411, USA. wmceneaney@ucsd.edu
Dept. of Industrial Eng. and Operations Research, Columbia University New York, NY, USA. antoinedesir@hotmail.com

payoff. The authors will attempt to ensure that the reader does not become confused between these two sub-games.

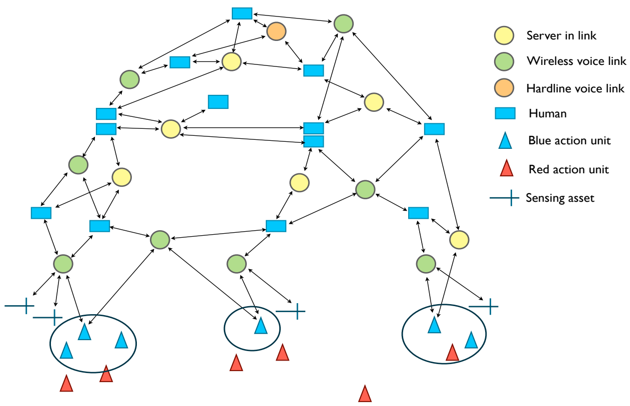

Fig. 1. Simple example network

III. Problem Definition

Prior to development of the problem model, we need to define the relevant mathematical objects. We first introduce the relevant idempotent algebras. The min-plus algebra (i.e., semifield) is given by

$$a \oplus b \doteq \min\{a, b\}, \qquad a \otimes b \doteq a + b,$$

operating on $I\!R^+ \doteq I\!R \cup \{+\infty\}$. The max-plus algebra (i.e., semifield) is given by

$$a \oplus^\vee b \doteq \min\{a, b\}, \qquad a \otimes^\vee b \doteq a + b,$$

operating on $I\!R^- \doteq I\!R \cup \{-\infty\}$. In the max-min algebra (i.e., semiring), the operations are defined as

$$a \vee b \doteq \max\{a, b\}, \qquad a \wedge b \doteq \min\{a, b\},$$

operating on $\overline{I\!R} \doteq I\!R \cup \{-\infty\} \cup \{+\infty\}$, where we note that $-\infty \vee b = b$ for all $b \in \overline{I\!R}$ and $-\infty \wedge b = -\infty$ for all $b \in \overline{I\!R}$ (c.f., [6]).

A. Payoff origin

We briefly describe the physical level game which will yield the payoff for the network disruption game that is our main focus. This will be a zero-sum game. Consider time t_k with $k \in]0, K[\doteq \{0, 1, 2 \ldots K\}$, where we note that throughout we will use the notation $]a, b[$ to denote $\{a, a+1, a+2 \ldots b\}$ for integers $a \leq b$. Suppose that at this time, there is an interaction between the Blue and Red action nodes. The possible actions for Blue are denoted by $v \in \mathcal{V} =]1, V[$. The true Red action node configuration is $x \in \mathcal{X} =]1, X[$. Given true asset configuration, x, Blue would obtain an action payoff $c(x, v)$. Let $C(v)$ be the vector of length X with components $c(x, v)$.

It is natural to use the max-plus probability structure (c.f., [1], [4], [8], [12] and the references therein) for deterministic games. (Note that, as with the networks, we model only the effect of Blue partial information.) Suppose Blue's knowledge of the true configuration is described by max-plus probability distribution, $q \in S^{\oplus^\vee X}$, where

$$S^{\oplus^\vee X} \doteq \Big\{ q = \in [-\infty, 0]^X \,\Big|\, \bigoplus_{x \in \mathcal{X}}{}^\vee q_x = 0 \Big\},$$

where $[-\infty, 0]$ denotes $(-\infty, 0] \cup \{-\infty\}$ and the X superscript indicates outer product X times. (Also, the $\bigoplus^\vee$ symbol indicates max-plus summation.) It is useful to recall that $[-\infty, 0]$ is analogous to $[0, 1]$ in the standard algebra, where 0 takes the place of 1 (the multiplicative identity), and $-\infty$ takes the place of 0 (the additive identity). We may interpret each component, q_x, as the additive inverse of the (relative) cost to Red to cause Blue to believe that the configuration is x. This will be become more clear below. The expected payoff for action $v \in \mathcal{V}$ given max-plus distribution q at time t_K, is as follows. Letting max-plus random variable ξ be distributed according to q, and $\mathbf{E}_q^{\oplus^\vee}$ denote min-plus expectation according to this q, the expected payoff is

$$\hat{J}(q, v) = \mathbf{E}_q^{\oplus^\vee}[c(v, \xi)] = \bigoplus_{x \in \mathcal{X}}{}^\vee c(v, x) \otimes^\vee q_x = C(v) \odot^\vee q,$$

where $\odot^\vee$ denotes the max-plus dot product. Given that Blue wants to minimize (make more negative) the physical level payoff, the value of information q is

$$\phi(q) \doteq \min_{v \in \mathcal{V}} J(q, v) = \bigwedge_{v \in \mathcal{V}} [C(v) \odot^\vee q]. \tag{1}$$

We see that if information is represented by a max-plus probability distribution over a finite set (and one has a finite set of controls), then *the value of information takes the form of a min-max sum of max-plus linear functionals over a max-plus probability simplex.* We also emphasize again that the above game merely provides the payoff for the network disruption game to follow.

B. Network game dynamics

It may again be helpful to refer to figure 1. We suppose that the Blue network will be defined as a graph, $(\mathcal{G}, \mathcal{E})$, where $\mathcal{G}$ denotes the set of nodes, and $\mathcal{E}$ denotes the set of edges. The set of nodes will be decomposed as $\mathcal{G} = \mathcal{G}_s \cup \mathcal{G}_a \cup \mathcal{G}_c \cup \mathcal{G}_d$ where $\mathcal{G}_s$ denotes the set of sensing nodes (air vehicle icons in the example figure), $\mathcal{G}_a$ denotes the set of action nodes (blue triangular icons in the example figure), $\mathcal{G}_c$ denotes the set of communication nodes (colored disc icons in the example figure) and $\mathcal{G}_d$ denotes the set of decision/analysis nodes (blue rectangular icons in the example figure). Let $\mathcal{G} =]1, n[$, and $\mathcal{E} = \{(g_i^1, g_i^2) \,|\, g_i^1, g_i^2 \in \mathcal{G} \, \forall i \in]1, n_e[\}$, where the elements of $\mathcal{E}$ are unordered pairs. Let $G_s = \#\mathcal{G}_s$, $G_a = \#\mathcal{G}_a$, $G_c = \#\mathcal{G}_c$ and $G_d = \#\mathcal{G}_d$. We suppose that for each action node, say $\alpha \in \mathcal{G}_a$, there exists a set of relevant sensing nodes, $\hat{\mathcal{G}}_s(\alpha) \subseteq \mathcal{G}_s$ such that information from these sensing elements affects the min-plus probability distribution describing information relevant to action node α. Further, given $\alpha \in \mathcal{G}_a$ and $\sigma \in \hat{\mathcal{G}}_s(\alpha)$, there exists an ordered set of decision/analysis nodes that information from σ must pass through prior to use by α. Let this path be denoted as

$$I^{\sigma,\alpha} = (\sigma, g_2, g_3 \dots g_{\bar{n}(\sigma,\alpha)-1}, \alpha)$$

with $g_k \in \mathcal{G}_d$ for all $k \in]2, \bar{n}(\sigma,\alpha) - 1[$. We also let $\mathcal{P}_N^{\sigma,\alpha}$ denote the set of ordered sequences of maximum length N given by

$$\mathcal{P}_N^{\sigma,\alpha} = \Big\{ \{\gamma_i\}_{i=1}^{\hat{n}} \,\Big|\, \hat{n} \in]1, N[, \, (\gamma_i, \gamma_{i+1}) \in \mathcal{E} \, \forall i \in]1, \hat{n}-1[, \text{ and s.t. } \exists \text{ subsequence } \{\gamma_{i_j}\}_{j=1}^{\bar{n}(\sigma,\alpha)} \text{ s.t. } i_1 = 1, \; i_{\bar{n}(\sigma,\alpha)} = \hat{n}, \, \gamma_{i_j} = g_j \in I^{\sigma,\alpha} \, \forall j \in]1, \bar{n}(\sigma,\alpha)[\Big\}.$$

Note that the elements of $\mathcal{P}_N^{\sigma,\alpha}$ are feasible paths through the graph passing through the required nodes of $I^{\sigma,\alpha}$ in the required order.

We now begin the discussion of the network disruption game. We use a fixed time-step model where, without loss of generality, $t_{k+1} - t_k = 1$ for all k. Let the delay in transfer of information along edge $(\gamma, g) \in \mathcal{E}$ be denoted by $\delta_{\gamma,g}^e \in [0,\infty)$, and the processing delay at node $\gamma \in \mathcal{G}$ be $\delta_\gamma^p \in [0,\infty)$. We let the delay at time t_k in arrival of information originating from node $\sigma \in \mathcal{G}_s$ at node $g \in \mathcal{G}$ be denoted by $d_{g,k}^\sigma$. We also let the vector of length n of such delays be d_k^σ. The delays in information arrival originating at all nodes in $\mathcal{G}_s$ will be $D_k = \{d_k^\sigma \,|\, \sigma \in \mathcal{G}_s\}$. Notationally, it will be helpful to arrange D_k as a vector of length nG_s. Of course, there may be multiple routes from node σ to node g. For $g \in \mathcal{G}$, let $\mathcal{N}_g \doteq \{\gamma \in \mathcal{G} \,|\, (\gamma, g) \in \mathcal{E}\}$. It is easy to see that the dynamics of $d_{g,k}^\sigma$ are given by

$$d_{g,k+1}^\sigma = \bigwedge_{\gamma \in \mathcal{N}_g} \bar{\delta}_{\gamma,g} + d_{g,k}^\sigma,$$

where $\bar{\delta}_{\gamma,g} = \delta_{\gamma,g}^e + \delta_\gamma^p$,

$$= \bigoplus_{\gamma \in \mathcal{N}_g} \bar{\delta}_{\gamma,g} \otimes d_{g,k}^\sigma, \tag{2}$$

where the $\bigoplus$ symbol denotes min-plus summation. Let T be the $n \times n$ matrix with elements

$$T_{g,\gamma} = \begin{cases} \bar{\delta}_{\gamma,g} & \text{if } \gamma \in \mathcal{N}_g, \\ +\infty & \text{otherwise.} \end{cases}$$

Then,

$$d_{g,k+1}^\sigma = \bigoplus_{\gamma \in \mathcal{G}} T_{g,\gamma} \otimes d_{\gamma,k}^\sigma \quad \forall \sigma \in \mathcal{G}_s, \, g \in \mathcal{G},$$

and arranging this in vector form,

$$d_{k+1}^\sigma = T \otimes d_k^\sigma \quad \forall \sigma \in \mathcal{G}_s, \tag{3}$$

where here $\otimes$ indicates min-plus matrix-vector multiplication. Also, letting $\overline{T}$ be the block-diagonal $nG_s \times nG_s$ matrix consisting of T blocks, with $+\infty$ (the min-plus zero) elsewhere, we may write

$$D_{k+1} = \overline{T} \otimes D_k. \tag{4}$$

Now suppose that Red may act to increase delays, and Blue may act to counter this. For example, in the case of radio transmissions, short delays may be induced by jamming, and with inclusion of intervening routers, delays may also be induced by other means. Finally, we arrive at

$$d_{k+1}^\sigma = T(u_k^b, u_k^r) \otimes d_k^\sigma \quad \forall \sigma \in \mathcal{G}_s, \tag{5}$$

or, equivalently,

$$D_{k+1} = \overline{T}(u_k^b, u_k^r) \otimes D_k, \tag{6}$$

where the Blue and Red controls at time t_k are $u_k^b \in \mathcal{U}^b$ and $u_k^r \in \mathcal{U}^r$, respectively, and where we assume $U^b = \#\mathcal{U}^b < \infty$ and $U^r = \#\mathcal{U}^r < \infty$. Note that (5), or equivalently (6), will define the dynamics in the network disruption game. *Importantly, we have a system with controlled (min-plus) linear dynamics.*

C. Network game payoff and value

We now begin definition of the payoff for this game. The payoff will be based on the information value given in (1), where we now need to determine the effect of delay on this quantity. The max-plus probability distribution, q, in (1) propagates in a surprisingly similar fashion to standard-sense probability distributions [8], [9].

Recall from Section III-A that we model Blue information states via some $q \in S^{\oplus^\vee X}$. Let us now add some more structure to this general form. We suppose the state can be decomposed according to domains partitioned by the action nodes. That is, we suppose the physical state, $x \in \mathcal{X}$, can be decomposed as $x = (x^1, x^2, \dots x^{G_a})$, with $x^a \in \mathcal{X}^a$ for all $a \in \mathcal{G}_a$, and $\mathcal{X} = \mathcal{X}^1 \times \mathcal{X}^2 \times \cdots \times \mathcal{X}^{G_a}$. With this decomposition, we may also let $q^a \in S^{\oplus^\vee X^a}$ where $X^a = \#\mathcal{X}^a$ for all $a \in \mathcal{G}_a$. For any $x \in \mathcal{X}$, $x = (x^1, x^2, \dots x^{G_a})$, we have $q_x = \bigotimes_{a \in \mathcal{G}_a}^\vee q_{x^a}^a$, and we note that by the usual summation process, one still has $q \in S^{\oplus^\vee X}$. This decomposition will be helpful in determining the cost of delay.

As discussed in [8], [9], we may suppose that in the absence of observations, information state, q_k, propoagates as a max-plus Markov chain. That is,

$$q_{k+1} = I\!\!P^T \otimes^\vee q_k,$$

where $I\!\!P$ is a max-plus probability transition matrix. In particular, $I\!\!P_{i,j} \in [-\infty, 0]$ for all $i, j \in]1, X[$ and $\bigoplus_{j=1}^{\vee X} I\!\!P_{i,j} = 0$.

Observation processing yields an update formula that is analogous to (standard-sense) Bayes rule (c.f., [8], [9]). We briefly recall how this occurs. Suppose q_k^a is the max-plus probability distribution of state component x^a at time t_k, prior to observation. Suppose that Blue obtains observation $y \in \mathcal{Y}$ (which we recall may be at least partially controlled by Red). Here, in order to reduce notation, we do not include the possible dependence of $\mathcal{Y}$ on domain. The resulting cost for any true state $x^a \in \mathcal{X}^a$ would be

$$\hat{q}_{x^a,k}^a = p^{\oplus^\vee}(y|x^a) + q_{x^a,k}^a = p^{\oplus^\vee}(y|x^a) \otimes^\vee q_{x^a,k}^a.$$

In solving the optimization problem, we are concerned only with the relative costs, and so we may normalize so that the max-plus sum over $x \in \mathcal{X}$ is zero. Let $\hat{q}_k^a$ (and q_k^a)

denote normalized costs, where we want $\bigoplus^{\vee}_{x^a \in \mathcal{X}^a} \hat{q}_{x^a,k} = \bigoplus^{\vee}_{x^a \in \mathcal{X}^a} q_{x^a,k} = 0$. The normalized cost is

$$\begin{aligned}\hat{q}^a_{x^a,k} &= p^{\oplus^\vee}(y|x^a) \otimes^\vee q^a_{x^a,k} - \Big\{ \bigoplus^{\vee}_{\zeta \in \mathcal{X}} \Big[p^{\oplus^\vee}(y|\zeta^a) \otimes^\vee q^a_{\zeta^a,k} \Big] \Big\} \\ &= p^{\oplus^\vee}(y|x^a) \otimes^\vee q^a_{x^a,k} \oslash^\vee \Big\{ \bigoplus^{\vee}_{\zeta^a \in \mathcal{X}^a} \Big[p^{\oplus^\vee}(y|\zeta^a) \otimes^\vee q^a_{\zeta^a,k} \Big] \Big\}, \quad (7)\end{aligned}$$

where $\oslash^\vee$ indicates max-plus division (standard-sense subtraction). We may interpret each component of the resulting max-plus probability at time t_k, $q_{x,k}$, as the additive inverse of the minimal relative cost to Red for modification of the observation process to yield observed sequence $\{y_0, y_1, \ldots y_k\}$ given true state x. We may denote the max-plus Bayes rule update as

$$\hat{q}^a_k = \mathcal{B}^y[q^a_k]. \quad (8)$$

Fix an action domain, indexed by $a \in \mathcal{G}^a$. For $y \in \mathcal{Y}$, let C^y be the $X^a \times X^a$ diagonal matrix with diagonal elements $p^{\oplus^\vee}(y|x^a)$. Also let R^y be the X^a-length vector with elements $p^{\oplus^\vee}(y|x^a)$. Written in vector form, update (7) takes the form

$$\hat{q}^a_k = \mathcal{B}^y[q^a_k] = [C^y \otimes^\vee q^a_k] \oslash^\vee [(R^y) \odot^\vee q^a_k]. \quad (9)$$

Now, suppose that there are multiple observation sources, indexed by $\sigma \in \hat{\mathcal{G}}_s(a)$. We would then have multiple observation processes for each domain, and we could denote the corresponding updates as

$$\hat{q}^a_k = \mathcal{B}^y_\sigma[q^a_k] = [C^y_\sigma \otimes^\vee q^a_k] \oslash^\vee [(R^y_\sigma) \odot^\vee q^a_k], \quad (10)$$

where the σ subscripts indicate the appropriate conditional probabilities corresponding to sensing node $\sigma \in \hat{\mathcal{G}}_s(a)$. Lastly, by padding these matrices and vectors appropriately, we may extend the updates to updates for the entire distribution, $q \in S^{\oplus^\vee X}$. That is, we may write the full update by sensing node $\sigma \in \mathcal{G}_s$ as

$$\hat{q}_k = \bar{\mathcal{B}}^y_\sigma[q_k] = [\bar{C}^y_\sigma \otimes^\vee q_k] \oslash^\vee [(\bar{R}^y_\sigma) \odot^\vee q_k], \quad (11)$$

For simplicity, let us suppose that each sensing node produces a single observation during each time-step. Then, the nominal dynamics of $q_\cdot$ take the form

$$q_{k+1} = I\!P^T \otimes^\vee \Big(\prod_{\sigma \in \mathcal{G}_s} \bar{\mathcal{B}}^y_\sigma \Big)[q_k], \quad (12)$$

where the $\prod$ notation indicates operator composition.

Now, recall again from Section III-A that we suppose Blue has some possible set of actions, possibly in response to Red actions, indexed by $v \in \mathcal{V}$. We assume Blue would like to minimize the max-plus expected cost given distribution $q \in S^{\oplus^\vee X}$. That is, the resulting payoff, assuming optimal Blue actions takes the form (as in Section III-A)

$$\bigwedge_{v \in \mathcal{V}} \mathbf{E}^{\oplus^\vee}_q[c(v,\xi)] = \bigwedge_{v \in \mathcal{V}} C(v) \odot^\vee q.$$

We now examine how this depends on network delays. The total expected payoff at time k of action at time $k+1$ given current distribution q_k then becomes

$$\psi^0(q_k) = \mathbf{E}^{\oplus^\vee} \Big\{ \bigwedge_{v \in \mathcal{V}} C(v) \odot^\vee q_{k+1} \Big\}, \quad (13)$$

where the expectation must be taken over the incoming observations. For notational simplicity, assume for the moment that there is only one sensor, σ. Substituting (11) and (12) into (13) yields

$$\begin{aligned}\psi^0(q_k) &= \bigoplus^{\vee}_{y \in \mathcal{Y}} \Big\{ \Big[\bigwedge_{v \in \mathcal{V}} C(v) \odot^\vee I\!P^T \otimes^\vee [\bar{C}^y_\sigma \otimes^\vee q_k] \\ &\qquad \oslash^\vee [(\bar{R}^y_\sigma) \odot^\vee q_k] \Big] \otimes^\vee \bigoplus^{\vee}_{\zeta \in \mathcal{X}} p^{\oplus^\vee}(y|\zeta) \otimes^\vee q_{\zeta,k} \Big\} \\ &= \bigoplus^{\vee}_{y \in \mathcal{Y}} \Big\{ \Big[\bigwedge_{v \in \mathcal{V}} C(v) \odot^\vee I\!P^T \otimes^\vee [\bar{C}^y_\sigma \otimes^\vee q_k] \\ &\qquad \oslash^\vee [(\bar{R}^y_\sigma) \odot^\vee q_k] \Big] \otimes^\vee (\bar{R}^y_\sigma) \odot^\vee q_k \Big\} \\ &= \bigoplus^{\vee}_{y \in \mathcal{Y}} \Big\{ \bigwedge_{v \in \mathcal{V}} \Big(\bar{C}^y_\sigma \otimes^\vee I\!P \otimes^\vee C(v) \Big) \odot^\vee q_k \Big\}.\end{aligned}$$

Now, returning to the case of multiple sensor nodes, and continuing to assume the same observation set for each, we obtain payoff form

$$\psi^0(q_k) = \bigoplus^{\vee}_{\vec{y} \in \mathcal{Y}^{G_s}} \Big\{ \bigwedge_{v \in \mathcal{V}} \Big[\Big(\bigotimes^{\vee}_{\sigma \in \mathcal{G}_s} \bar{C}^{y_\sigma}_\sigma \Big) \otimes^\vee I\!P \otimes^\vee C(v) \Big] \odot^\vee q_k \Big\}. \quad (14)$$

Next, we model the effects of delays through this payoff model. For the moment, we suppose these delays are all discretized to the same time-step as the dynamics, where k indexes time-step t_k. If information from node $\sigma \in \hat{\mathcal{G}}_s(\alpha)$ is delayed at $\alpha \in \mathcal{G}_a$ by some time, d^σ_α, then the observation updates in (14) for time-steps more recent than d^σ_α time back, will not take place. Suppose that at time k, the maximum delay is $d^*_k = \max_{\alpha \in \mathcal{G}_a} \max_{\sigma \in \hat{\mathcal{G}}_s(\alpha)} d^\sigma_{\alpha,k}$. Let $j^* = j^*(k) \doteq \lfloor (k - d^*_k)$. For simplicity, we assume $\hat{\mathcal{G}}_s(\alpha_1) \cap \hat{\mathcal{G}}_s(\alpha_2) = \emptyset$ if $\alpha_1 \neq \alpha_2$. For each $j \in]j^*(k), k[$, and each $\alpha \in \mathcal{G}_a$, let $\tilde{\mathcal{G}}_s(\alpha, j, k) \doteq \{\sigma \in \hat{\mathcal{G}}_s(\alpha) \,|\, d^\sigma_{\alpha,k} < k - j\}$. Also let $\bar{\mathcal{G}}_s(j,k) \doteq \bigcup_{\alpha \in \mathcal{G}_a} \tilde{G}_s(\alpha, j, k)$ and $\bar{G}_s(j,k) = \#\bar{\mathcal{G}}_s(j,k)$. For $\vec{y} \in \mathcal{Y}^{\bar{G}_s(j,k)}$ indexed by $\sigma \in \bar{\mathcal{G}}_s(j,k)$, let $\widetilde{C}^{\vec{y}} \doteq \bigotimes^{\vee}_{\sigma \in \bar{\mathcal{G}}_s(j,k)} \bar{C}^{y_\sigma}_\sigma$. Next, let $\hat{\mathcal{Y}}_k \doteq \mathcal{Y}^{\bar{G}_s(j^*,k)} \times \mathcal{Y}^{\bar{G}_s(j^*+1,k)} \times \cdots \times \mathcal{Y}^{\bar{G}_s(k,k)}$. For $\vec{\vec{y}} \in \hat{\mathcal{Y}}_k$, let

$$\mathcal{O}^{\vec{\vec{y}}} \doteq \bigotimes^{\vee}_{j \in]j^*,k[} \Big\{ \Big[\bigotimes^{\vee}_{\sigma \in \bar{\mathcal{G}}_s(j,k)} \bar{C}^{y_\sigma}_\sigma \Big] \otimes^\vee I\!P \Big\}.$$

For the purposes of measuring the cost of delay, we suppose that the information state system is in some state, $\bar{q}$, at the initial time. The cost of delays D_k is given by

$$\begin{aligned}\psi &= \bigoplus^{\vee}_{\vec{\vec{y}} \in \hat{\mathcal{Y}}_k} \Big\{ \bigwedge_{v \in \mathcal{V}} \Big[\mathcal{O}^{\vec{\vec{y}}} \otimes^\vee C(v) \Big] \odot^\vee \bar{q} \Big\} \\ &= \bigvee_{\vec{\vec{y}} \in \hat{\mathcal{Y}}_k} \Big\{ \bigwedge_{v \in \mathcal{V}} \Big[\mathcal{O}^{\vec{\vec{y}}} \otimes^\vee C(v) \Big] \odot^\vee \bar{q} \Big\},\end{aligned}$$

where we are intentionally leaving the arguments of ψ unspecified for the moment. Now, using the min-max distributive property, this becomes

$$\psi = \bigwedge_{\{v_{\vec{\vec{y}}}\}\in\hat{\mathcal{V}}} \left\{ \bigvee_{\vec{\vec{y}}\in\hat{\mathcal{Y}}_k} \left[\mathcal{O}^{\vec{\vec{y}}} \otimes^{\vee} C(v) \right] \odot^{\vee} \bar{q} \right\}$$

$$= \bigwedge_{\{v_{\vec{\vec{y}}}\}\in\hat{\mathcal{V}}} \left\{ \bigoplus_{\vec{\vec{y}}\in\hat{\mathcal{Y}}_k}^{\vee} \left[\mathcal{O}^{\vec{\vec{y}}} \otimes^{\vee} C(v) \right] \odot^{\vee} \bar{q} \right\},$$

where $\hat{\mathcal{V}}$ denotes the set of sequences of length $\#\hat{\mathcal{Y}}_k$ of elements of $\mathcal{V}$ indexed by $\vec{\vec{y}} \in \mathcal{Y}_k$. Substituting back in for $\mathcal{O}^{\vec{\vec{y}}}$, we have

$$\psi = \bigwedge_{\{v_{\vec{\vec{y}}}\}\in\hat{\mathcal{V}}} \left\{ \bigoplus_{\vec{\vec{y}}\in\hat{\mathcal{Y}}_k}^{\vee} \left[\bigotimes_{j\in]j^*,k[}^{\vee} \left[\bigotimes_{\sigma\in\bar{\mathcal{G}}_s(j,k)}^{\vee} \bar{C}_{\sigma}^{y_{\sigma,j}} \right] \otimes^{\vee} \mathcal{L}(v_{\vec{\vec{y}}}) \right] \odot^{\vee} \bar{q} \right\},$$

where $\mathcal{L}(v) \doteq I\!\!P \otimes^{\vee} C(v)$. Given that we are assuming $\hat{\mathcal{G}}_s(\alpha_1) \cap \hat{\mathcal{G}}_s(\alpha_2) = \emptyset$ if $\alpha_1 \neq \alpha_2$, we may simply let $\hat{\alpha}(\sigma)$ be the action node affected by sensor node σ. Reordering the products, one obtains

$$\psi = \bigwedge_{\{v_{\vec{\vec{y}}}\}\in\hat{\mathcal{V}}} \left\{ \bigoplus_{\vec{\vec{y}}\in\hat{\mathcal{Y}}_k}^{\vee} \left[\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \left[\bigotimes_{j\in]1,d^{\sigma}_{\hat{\alpha}(\sigma),k}[}^{\vee} \bar{C}_{\sigma}^{y_{\sigma,j}} \right] \otimes^{\vee} \mathcal{L}(v_{\vec{\vec{y}}}) \right] \odot^{\vee} \bar{q} \right\},$$

which upon letting $\delta \in \mathbf{N}^{G_s}$ generically denote $d^{\cdot}_{\hat{\alpha}(\cdot),k}$,

$$= \bigwedge_{\{v_{\vec{\vec{y}}}\}\in\hat{\mathcal{V}}} \left\{ \bigoplus_{\vec{\vec{y}}\in\hat{\mathcal{Y}}_k}^{\vee} \left[\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \left[\bigotimes_{j\in]1,\delta_\sigma[}^{\vee} \bar{C}_{\sigma}^{y_{\sigma,j}} \right] \otimes^{\vee} \mathcal{L}(v_{\vec{\vec{y}}}) \right] \odot^{\vee} \bar{q} \right\}$$

$$= \bigwedge_{\{v_{\vec{\vec{y}}}\}\in\hat{\mathcal{V}}} \left\{ \bigoplus_{\vec{y}\in\mathcal{Y}^{\delta_{\sigma_1}}}^{\vee} \bigoplus_{\vec{y}\in\mathcal{Y}^{\delta_{\sigma_2}}}^{\vee} \cdots \bigoplus_{\vec{y}\in\mathcal{Y}^{\delta_{\sigma_{G_s}}}}^{\vee} \left[\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \left(\bigotimes_{j\in]1,\delta_\sigma[}^{\vee} \bar{C}_{\sigma}^{y_{\sigma,j}} \right) \right] \otimes^{\vee} \mathcal{L}(v_{\vec{\vec{y}}}) \right] \odot^{\vee} \bar{q} \right\},$$

which upon rearrangement and a little work to see that it is sufficient to consider only $\{v_\delta\} \in \tilde{\mathcal{V}}$, the set of sequences indexed by delay vectors δ,

$$= \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \left\{ \left[\left[\left[\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \left(\bigoplus_{\vec{y}\in\mathcal{Y}^{\delta_\sigma}}^{\vee} \bigotimes_{j\in]1,\delta_\sigma[}^{\vee} \bar{C}_{\sigma}^{y_{\sigma,j}} \right) \right] \otimes^{\vee} \mathcal{L}'(\{v_\delta\}) \right] \odot^{\vee} \bar{q} \right\}.$$

Letting, for $\delta_\sigma \in \mathbf{N}$, $\tilde{C}_{\sigma}^{\delta_\sigma} \doteq \bigoplus_{\vec{y}\in\mathcal{Y}^{\delta_\sigma}}^{\vee} \bigotimes_{j\in]1,\delta_\sigma[}^{\vee} \bar{C}_{\sigma}^{y_{\sigma,j}}$, we have

$$\psi(\delta) = \psi(\delta;\bar{q}) = \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \left\{ \left[\left[\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \tilde{C}_{\sigma}^{\delta_\sigma} \right] \otimes^{\vee} \mathcal{L}'(\{v_\delta\}) \right] \odot^{\vee} \bar{q} \right\}, \quad (15)$$

where we have now indicated the arguments of ψ, and we note that we are specifically interested in the dependence on our state variable, D_k, where we note that $\delta = d^{\cdot}_{\hat{\alpha}(\cdot),k}$ form a subset of the components of D_k.

Noting again that the $\tilde{C}_{\sigma}^{\delta_\sigma}$ are diagonal, (15) becomes

$$\psi(\delta) = \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \bigvee_{x\in\mathcal{X}} \left\{ \left(\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \tilde{C}_{\sigma}^{\delta_\sigma} \right)_{x,x} \otimes^{\vee} \mathcal{L}'_x(\{v_\delta\}) \otimes^{\vee} \bar{q}_x \right\}.$$

Also, noting that we are interested in the cost relative to no delay, we let

$$\hat{\psi}(\delta) \doteq \psi(\delta) - \psi(0) = \psi(\delta;\bar{q}) - \psi(0;\bar{q})$$

$$= \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \bigvee_{x\in\mathcal{X}} \left\{ \left(\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \tilde{C}_{\sigma}^{\delta_\sigma} \right)_{x,x} \otimes^{\vee} \mathcal{L}'_x(\{v_\delta\}) \otimes^{\vee} \bar{q}_x \otimes^{\vee} (-\psi(0)) \right\}$$

$$\doteq \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \bigvee_{x\in\mathcal{X}} \left\{ \left(\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \tilde{C}_{\sigma}^{\delta_\sigma} \right)_{x,x} \otimes^{\vee} \mathcal{L}'_x(\{v_\delta\}) \otimes^{\vee} \bar{q}'_x \right\},$$

and letting $\hat{\mathcal{L}}_x(\{v_\delta\}) \doteq \mathcal{L}'_x(\{v_\delta\}) \otimes^{\vee} \bar{q}'_x$ for all $x \in \mathcal{X}$,

$$= \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \bigvee_{x\in\mathcal{X}} \left\{ \left(\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \tilde{C}_{\sigma}^{\delta_\sigma} \right)_{x,x} \otimes^{\vee} \hat{\mathcal{L}}_x(\{v_\delta\}) \right\}. \quad (16)$$

Note that $\hat{\psi}$ will be zero at the origin, monotonically increasing with respect to the partial order, Lipschitz, and bounded above by the maximum loss cost. Consequently (c.f., [5]), and noting that we have only specified the data at discrete values of δ, without loss of generality, we may write it in the form of a finite-complexity, min-plus convex function over the first octant in $I\!\!R^{G_s}$. Let

$$\left(\bigotimes_{\sigma\in\mathcal{G}_s}^{\vee} \tilde{C}_{\sigma}^{\delta_\sigma} \right)_{x,x} = \bigvee_{z\in\mathcal{Z}} \left[a^{z,x} \oplus \bigoplus_{\sigma\in\mathcal{G}_s} b_{\sigma}^{z,x} \otimes c_\sigma \delta_\sigma \right],$$

for some finite index set, $\mathcal{Z}$. The c_σ (standard-sense) multipliers are not strictly necessary, but appear to be helpful in computational-complexity reduction. By the change of variables, $\bar{\delta}_\sigma = c_\sigma \delta_\sigma$ for all $\sigma \in \mathcal{G}_s$, we may, with a bit more work, remove the multipliers. With an abuse of notation, we replace $\bar{\delta}$ with δ throughout. We have

$$\hat{\psi}(\delta) = \bigwedge_{\{v_\delta\}\in\tilde{\mathcal{V}}} \bigvee_{x\in\mathcal{X}} \left(\bigvee_{z\in\mathcal{Z}} \left[a^{z,x} \oplus \bigoplus_{\sigma\in\mathcal{G}_s} b_{\sigma}^{z,x} \otimes \delta_\sigma \right] \right) \otimes \hat{\mathcal{L}}_x(\{v_\delta\}),$$

and reducing notation by letting the elements of $\tilde{\mathcal{V}}$ be denoted simply as ν,

$$= \bigwedge_{\nu\in\tilde{\mathcal{V}}} \bigvee_{x\in\mathcal{X}} \left(\bigvee_{z\in\mathcal{Z}} \left[a^{z,x} \oplus \bigoplus_{\sigma\in\mathcal{G}_s} b_{\sigma}^{z,x} \otimes \delta_\sigma \right] \right) \otimes \hat{\mathcal{L}}_x(\nu).$$

Let $\hat{\mathcal{Z}} = \mathcal{X} \times \mathcal{Z}$, and let the elements be denoted simply as $\hat{z}$. Also, let $a_1^{\hat{z},\nu} \doteq a^{z,x} \otimes \hat{\mathcal{L}}_x(\nu)$ and $b_{1,\sigma}^{\hat{z},\nu} \doteq b_{\sigma}^{z,x} \otimes \hat{\mathcal{L}}_x(\nu)$. We have

$$\hat{\psi}(\delta) = \bigwedge_{\nu\in\tilde{\mathcal{V}}} \bigvee_{\hat{z}\in\hat{\mathcal{Z}}} \left[a_1^{\hat{z},\nu} \oplus \bigoplus_{\sigma\in\mathcal{G}_s} b_{1,\sigma}^{\hat{z},\nu} \otimes \delta_\sigma \right]$$

$$= \bigwedge_{\nu\in\tilde{\mathcal{V}}} \bigvee_{\hat{z}\in\hat{\mathcal{Z}}} \left[a_1^{\hat{z},\nu} \oplus b_1^{\hat{z},\nu} \odot \delta \right].$$

Applying the max-min distributive property, and letting $\tilde{Z}$ (with elements denoted as $\tilde{z}$) be the set of sequences of elements of $\hat{Z}$ indexed by $\nu \in \tilde{\mathcal{V}}$, this yields (where we also note $\wedge \equiv \oplus$)

$$\hat{\psi}(\delta) = \bigvee_{\tilde{z}\in\tilde{\mathcal{Z}}} \bigoplus_{\nu\in\tilde{\mathcal{V}}} \left[a_1^{\tilde{z}_\nu,\nu} \oplus b_1^{\tilde{z}_\nu,\nu} \odot \delta\right] \doteq \bigvee_{\tilde{z}\in\tilde{\mathcal{Z}}} \left[a_2^{\tilde{z}} \oplus b_2^{\tilde{z}} \odot \delta\right], \tag{17}$$

where $a_2^{\tilde{z}} = \bigoplus_{\nu\in\tilde{\mathcal{V}}} a_1^{\tilde{z}_\nu,\nu}$ and $b_2^{\tilde{z}} = \bigoplus_{\nu\in\tilde{\mathcal{V}}} b_1^{\tilde{z}_\nu,\nu}$. This is a finite-complexity min-plus convex function. Here, by finite complexity, we mean a min-plus convex function composed as a pointwise maximum of min-plus affine functions (i.e., a max-min linear combination of min-plus linear functions). Again, we remind the reader that the space of min-plus convex functions is a max-min linear space, in analogy with the space of standard-sense convex functions as a max-plus linear space [5].

The function, $\hat{\psi}$ will be the running cost in our delay game. We suppose a finite time-horizon problem formulation, with terminal time K (i.e., t_K). We also take a maximum over the time-domain. With this, the payoff for our game with initial time, k_0, and initial state, D_{k_0}, is given by

$$\bar{J}(k_0, D_{k_0}, u_\cdot^b, u_\cdot^r) = \bigvee_{k\in]k_o,K[} \hat{\psi}(D_k), \tag{18}$$

for $u_\cdot^b \in (\mathcal{U}^b)^{K-k_0}$ and $u_\cdot^r \in (\mathcal{U}^r)^{K-k_0}$ where the superscript $K-k_0$ indicates outer product, and where we abuse notation by using the full D_k vector as the argument for $\hat{\psi}$, which only depends on the $d^\sigma_{\hat{\alpha}(\sigma),l}$ components. Of course, it is implicit that the dynamics for $D_\cdot$ are given by (4).

As we desire a risk-averse/worst-case analysis, we consider the upper value of the game. Let

$$\mathcal{R}_{k_0} \doteq \{\rho : (\mathcal{U}^b)^{K-k_0} \to (\mathcal{U}^r)^{K-k_0} \,|\, \text{nonanticipative}\}.$$

The upper value will be

$$\overline{W}(k_0, D_{k_0}) = \bigvee_{\rho\in\mathcal{R}_{k_0}} \bigwedge_{u_\cdot^b\in(\mathcal{U}^b)^{K-k_0}} \bar{J}(k_0, D_{k_0}, u_\cdot^b, \rho[u_\cdot^b]). \tag{19}$$

IV. Min-Plus Convex Form Propagation

We now have a payoff, $\bar{J}$, given as a maximum over time, i.e., as a max-plus sum over time, rather than as a standard-sense sum over time. The more-complex continuous-time case has been considered in, for example [4] among numerous sources. Here, the finite max-plus summation leads to a simpler analysis. We briefly discuss this analysis, yielding the dynamic program, as it impacts the rather interesting propagation form.

Consider an initial time $k \in]k_0, K[$. The value is given by

$$\overline{W}(k, \delta) = \bigvee_{\rho\in\mathcal{R}_k} \bigwedge_{u_\cdot^b\in(\mathcal{U}^b)^{K-k}} \bigvee_{j\in]k,K[} \hat{\psi}(D_j),$$

where $D_\cdot$ satisfies (4) with controls $u_\cdot^b$ and $\rho(u_\cdot^b)$ and initial state $D_k = \delta$. With a little work, one finds

$$\begin{aligned}\overline{W}(k,\delta) &= \bigwedge_{u_k^b\in\mathcal{U}^b} \bigvee_{u_k^r\in\mathcal{U}^r} \Big\{\hat{\psi}(D_k) \vee \Big[\\ &\qquad \bigvee_{\rho\in\mathcal{R}_{k+1}} \bigwedge_{u_\cdot^b\in(\mathcal{U}^b)^{K-(k+1)}} \bigvee_{j\in]k+1,K[} \hat{\psi}(D_j)\Big]\Big\} \\ &= \bigwedge_{u^b\in\mathcal{U}^b} \bigvee_{u^r\in\mathcal{U}^r} \Big\{\hat{\psi}(D_k) \vee \overline{W}(k+1, D_{k+1})\Big\} \\ &= \hat{\psi}(D_k) \vee \Big[\bigwedge_{u^b\in\mathcal{U}^b} \bigvee_{u^r\in\mathcal{U}^r} \overline{W}(k+1, D_{k+1})\Big],\end{aligned} \tag{20}$$

where $D_{k+1} = \overline{T}(u^b, u^r)\delta$.

Now suppose

$$\overline{W}(k+1, \delta) = \bigvee_{\omega\in\Omega_{k+1}} \Big[\bar{a}^\omega_{k+1} \oplus \bar{b}^\omega_{k+1} \odot \delta\Big], \tag{21}$$

which is certainly true at time t_K, where $\overline{W}(K, \delta) = \hat{\psi}(\delta)$. Then by (20) and (17),

$$\begin{aligned}\overline{W}(k,\delta) &= \Big[\bigvee_{\tilde{z}\in\tilde{\mathcal{Z}}} \big(a_2^{\tilde{z}} \oplus b_2^{\tilde{z}} \odot \delta\big)\Big] \\ &\qquad \vee \Big[\bigwedge_{u^b\in\mathcal{U}^b} \bigvee_{u^r\in\mathcal{U}^r} \bigvee_{\omega\in\Omega_{k+1}} \Big(\bar{a}^\omega_{k+1} \oplus \bar{b}^\omega_{k+1} \odot D_{k+1}\Big)\Big] \\ &= \Big[\bigvee_{\tilde{z}\in\tilde{\mathcal{Z}}} \big(a_2^{\tilde{z}} \oplus b_2^{\tilde{z}} \odot \delta\big)\Big] \\ &\vee \Big[\bigwedge_{u^b\in\mathcal{U}^b} \bigvee_{u^r\in\mathcal{U}^r} \bigvee_{\omega\in\Omega_{k+1}} \Big(\bar{a}^\omega_{k+1} \oplus (\overline{T}^T(u^b, u^r) \otimes \bar{b}^\omega_{k+1}) \odot \delta\Big)\Big],\end{aligned}$$

and letting $\hat{b}_k^{\omega,u^b,u^r} \doteq \overline{T}^T(u^b, u^r) \otimes \bar{b}^\omega_{k+1}$,

$$\begin{aligned}&= \Big[\bigvee_{\tilde{z}\in\tilde{\mathcal{Z}}} \big(a_2^{\tilde{z}} \oplus b_2^{\tilde{z}} \odot \delta\big)\Big] \\ &\quad \vee \Big[\bigwedge_{u^b\in\mathcal{U}^b} \bigvee_{(\omega,u^r)\in\Omega_{k+1}\times\mathcal{U}^r} \Big(\bar{a}^\omega_{k+1} \oplus \hat{b}_k^{\omega,u^b,u^r} \odot \delta\Big)\Big].\end{aligned} \tag{22}$$

The last term in square brackets in (22) is a max-min product of min-plus finite-complexity convex functions. We discuss this in more detail in Section V, but first, note that by applying the max-min distributive property to ereq:Wbarup1, one obtains

$$\begin{aligned}\overline{W}(k,\delta) &= \Big[\bigvee_{\tilde{z}\in\tilde{\mathcal{Z}}} \big(a_2^{\tilde{z}} \oplus b_2^{\tilde{z}} \odot \delta\big)\Big] \\ &\vee \Big[\bigvee_{\{(\omega,u^r)_{u^b}\}\in(\Omega_{k+1}\times\mathcal{U}^r)^{U^b}} \bigoplus_{u^b\in\mathcal{U}^b} \Big(\bar{a}^\omega_{k+1} \oplus \hat{b}_k^{\omega,u^b,u^r} \odot \delta\Big)\Big],\end{aligned}$$

and letting $\tilde{a}_k^{\{\omega_{u^b}\}} \doteq \bigoplus_{u^b\in\mathcal{U}^b} \bar{a}^{\omega_{u^b}}_{k+1}$ and $\tilde{b}_k^{\{(\omega,u^r)_{u^b}\}} \doteq \bigoplus_{u^b\in\mathcal{U}^b} \hat{b}_k^{(\omega,u^r)_{u^b},u^b}$, this becomes

$$= \left[\bigvee_{\tilde{z} \in \tilde{\mathcal{Z}}} \left(a_2^{\tilde{z}} \oplus b_2^{\tilde{z}} \odot \delta \right) \right]$$

$$\vee \left[\bigvee_{\{(\omega,u^r)_{ub}\} \in (\Omega_{k+1} \times \mathcal{U}^r)^{U^b}} \left(\tilde{a}_k^{\{\omega_{ub}\}} \oplus \tilde{b}_k^{\{(\omega,u^r)_{ub}\}} \odot \delta \right) \right]$$

$$= \bigvee_{\omega \in \Omega_k} \left[\bar{a}_k^\omega \oplus \bar{b}_k^\omega \odot \delta \right] = \bigvee_{\omega \in \Omega_k} \left[\bar{a}_k^\omega \wedge \bar{b}_k^\omega \odot \delta \right]. \quad (23)$$

for proper relabeling and reindexing. where in the last form, we emphasize that this is a max-min linear combination of min-plus linear functionals (i.e., a finite-complexity min-plus convex functional). We have

Theorem 4.1: For all $k \in]0, K[$, $\overline{W}$ takes the form (23). Further, one obtains $\{\bar{a}_k^\omega \,|\, \omega \in \Omega_k\}$ and $\{\bar{b}_k^\omega \,|\, \omega \in \Omega_k\}$ from $\{\bar{a}_{k+1}^\omega \,|\, \omega \in \Omega_{k+1}\}$ and $\{\bar{b}_{k+1}^\omega \,|\, \omega \in \Omega_{k+1}\}$ by idempotent linear algebraic operations.

In short, value function propagation reduces to idempotent linear operations.

V. Complexity Analysis

We see that the solution of our network disruption game reduces to propagation of coefficients by idempotent linear operations, where the value function is given as a max-min linear combination of min-plus linear functionals at each time-step. However, the apparent cardinality of the set of coefficients grows extremely rapidly as one back-propagates. In this section, we discuss the complexity of such forms and optimal max-min linear projections for complexity-growth attenuation.

A. Complexity of max-min sums and products of min-plus linear functions

From an examination of the value function propagation derived in the previous section, one sees that there appear to be a tremendous number of new min-plus affine functions added to the value-function representation at each time-step. In practice, the overwhelming majority of these play no role. Typically they are either duplicates of existing functions, or they are inactive, by which we mean that they nowhere achieve the pointwise maximum in (23). We say a component affine function is strictly active if it is the unique maximizing function at some point. Consequently, they can be eliminated from the representation with no loss of accuracy. Within the limited space here, let us give some short synopsis of the theory underlying max-min sums and products of min-plus affine functions.

First, we deal with finite-complexity min-plus convex functions. The general form is

$$f(d) = \bigvee_{j \in \mathcal{J}} \left[a^j \wedge b^j \odot d \right] = \bigvee_{j \in \mathcal{J}} \left[a^j \oplus b^j \odot d \right] \doteq \bigvee_{j \in \mathcal{J}} h^j(d), \quad (24)$$

where $\mathcal{J} =]1, J[$, $d \in (I\!R^+)^n$, $a^j \in I\!R^+$ and $b^j \in (I\!R^+)^n$ for all $j \in \mathcal{J}$. Let $\mathcal{N} =]1, n[$. (Here, for simplicity, we suppose the functions are defined over all of $(I\!R^+)^n$ rather than only $[0, +\infty]^n$.) An important object with respect to each of the min-plus affine components, $h^j(\cdot)$, is its crux. The crux value of h^j is $v^j \doteq a^j$. The crux location, c^j, is given componentwise as $c_i^j = v^j - b_i^j$ for all $i \in]1, n[$. The crux is the pair, (c^j, v^j). Intuitively speaking, the crux is the unique point where the $n+1$ hyperplane sections which form the graph of h^j intersect. The importance of the crux with regard to these expansions is evident from the following simple result. Note that we suppose that duplicates have already been removed from representation (24) (i.e., there do not exist $\bar{j} \neq \hat{j}$ such that $b^{\bar{j}} = b^{\hat{j}}$ and $a^{\bar{j}} = a^{\hat{j}}$).

Lemma 5.1: A min-plus affine functional is strictly active in (24) if and only if it is strictly active at its crux location.

Proof: Sufficiency is obvious, and so we only prove necessity. Suppose affine function $h^{\bar{j}}$ is strictly active. Then, there exists $\hat{d} \in I\!R^n$ such that

$$h^{\bar{j}}(\hat{d}) = a^{\bar{j}} \oplus b^{\bar{j}} \odot \hat{d} > a^j \oplus b^j \odot \hat{d} = h^j(\hat{d}) \quad \forall j \in \mathcal{J} \setminus \{\bar{j}\}. \quad (25)$$

Fix any $j \neq \bar{j}$.

Suppose $a^j \leq b^j \odot \hat{d}$. Then, by (25), $a^{\bar{j}} \oplus b^{\bar{j}} \odot \hat{d} > a^j$. This implies

$$h^{\bar{j}}(c^{\bar{j}}) = v^{\bar{j}} = a^{\bar{j}} > a^j \geq a^j \oplus b^j \odot c^{\bar{j}} = h^j(c^{\bar{j}}),$$

which is the desired result.

Now instead, suppose $a^j > b^j \odot \hat{d}$. Then, by (25), $a^{\bar{j}} \oplus b^{\bar{j}} \odot \hat{d} > b^j \odot \hat{d}$, which implies

$$b^{\bar{j}} \odot \hat{d} > b^j \odot \hat{d}. \quad (26)$$

Let $\tilde{i} \in \mathcal{N}$ be such that

$$b_{\tilde{i}}^j \otimes \hat{d}_{\tilde{i}} = b^j \odot \hat{d}. \quad (27)$$

By (26) and (27), $b_{\tilde{i}}^{\bar{j}} \otimes \hat{d}_{\tilde{i}} > b_{\tilde{i}}^j \otimes \hat{d}_{\tilde{i}}$, which implies $b_{\tilde{i}}^{\bar{j}} > b_{\tilde{i}}^j$, and consequently, $\bigoplus_{i \in \mathcal{N}} b_i^j - b_i^{\bar{j}} < 0$. This implies

$$h^{\bar{j}}(c^{\bar{j}}) = v^{\bar{j}} > \bigoplus_{i \in \mathcal{N}} b_i^j + v^{\bar{j}} - b_i^{\bar{j}} = \bigoplus_{i \in \mathcal{N}} b_i^j + c^{\bar{j}}$$
$$= b^j \odot c^{\bar{j}} \geq a^j \oplus b^j \odot c^{\bar{j}} = h^j(c^{\bar{j}}).$$

■

We remark that the cruxes in a finte-complexity min-plus convex function can also be defined geometrically, without recourse to a representation such as (24), and we do not include the details. One finds the following.

Lemma 5.2: The number of active affine functions in f is exactly the number of cruxes.

We say that a finite-complexity min-plus convex functional given as (24) is in minimal form if it does not contain any inactive affine functionals. We also note the following useful result, and do not include the straight-forward proof.

Theorem 5.3: Suppose finite-complexity min-plus convex function $f : (I\!R^+)^n \to I\!R^+$ has exactly J cruxes, and these are $\mathcal{C} \doteq \{(c^j, v^j) \,|\, j \in \mathcal{J}\}$ where $\mathcal{J} =]1, J[$. Let $\hat{f} : (I\!R^+)^n \to I\!R^+$ be given by $\hat{f}(d) = \bigvee_{j \in \mathcal{J}} a^j \otimes b^j \odot d$, where $a^j = v^j$ and $b_i^j = v^j - c_i^j$ for all $j \in \mathcal{J}$ and all $i \in \mathcal{N}$. Then, $f = \hat{f}$, and this is the unique minimal form.

Max-min sums of finite-complexity min-plus convex functions increase the complexity at most linearly, with this growth often being tempered by crux dominance. On the other hand, in the value function propagation derivation, we had max-min *products* of finite-complexity min-plus convex functions. Very interestingly, because of the max-min distributive property, these may continue to be represented as finite-complexity min-plus convex functions. However, the complexity appears to grow very rapidly: With a product over K groups of J-complexity min-plus convex functions, the apparent complexity, purely from examination of the distributive property, is the cardinality of sequences of length K of elements of J, i.e., J^K. However, there is another bound, induced by the geometry, that can be significantly lower. There is some similarity in this discussion to the standard-sense finite-complexity discussion in [14]. We very briefly describe the results. Suppose

$$g(d) = \bigwedge_{k\in\mathcal{K}} \bigvee_{j\in\mathcal{J}} h^{k,j}(d) = \bigwedge_{k\in\mathcal{K}} \bigvee_{j\in\mathcal{J}} \Big[a^{k,j} \oplus b^{k,j} \odot d\Big]. \quad (28)$$

In the one-dimensional case, $n = 1$, a very nice complexity bound is obtained. (Because of space limitations, we do not include a proof.)

Theorem 5.4: In the case of domain, $I\!R^+$, the complexity of (28) is at most $KJ - (K - 1)$, where this bound is tight.

Obviously, this can be much better than J^K. Tight bounds in higher dimensions remain an open question. However, one known bound is $(KJ)^{n-1}$, where again, space limitations do not allow for a proof.

B. Max-min projection and complexity attenuation

Regardless of complexity bounds, there is still a need for complexity-growth attenuation. For the general class of finite-complexity min-plus convex functions, we extend the results of [5]. There are two questions. First, what is the optimal (minimum error) complexity-reduction representation? Second, how does one compute this representation? Due to space limitations, we only briefly indicate the results.

In regard to the first question, results for finite-complexity standard-sense convex functions were given in [10]. These results were extended to finite-complexity max-plus convex functions in [5]. In the case at hand, one trivially transfers the results for the max-plus case to the min-plus case. Suppose one has $f(d) = \bigvee_{j\in\mathcal{J}} h^j(d)$ with $J = \#\mathcal{J}$, and wishes to find an approximation, $g(d) = \bigvee_{m\in\mathcal{M}} \hat{h}^m(d)$ with $M = \#\mathcal{M} < J$. Note that as this is a pointwise maximum, any error such that $g(\tilde{d}) > f(\tilde{d})$ for some $\tilde{d} \in I\!R^n$ cannot be corrected by the addition of more terms to g. Consequently, one seeks an approximation from below.

For given $x \in I\!R^k$ for some $k \in \mathbf{N}$, w define the downward cone as $\mathcal{D}(x) = \{\hat{x} \in I\!R^k \,|\, \hat{x} \preceq x\}$ where $\preceq$ denotes the partial order on $I\!R^k$. Next, for a set of points, $X \subset I\!R^k$, we let $\langle X\rangle$ denote the convex hull of X. Then, we may define the *min-plus cornice* of X as

$$\mathcal{C}^{\oplus}(X) = \bigcup_{x\in\langle X\rangle} \mathcal{D}(x).$$

One may show that our optimal complexity reduction problem reduces to maximization of a min-plus convex, monotonically increasing function over an outer-product of cornices, where the cornices are formed from the coefficients describing f. We find, that the optimal solution (also the optimal min-plus projection) is obtained by pruning the set of constituent min-plus affine functions describing f down to M elements (see [5]).

The second question regards the computational cost of this optimal projection/pruning. Noting the discussion of cruxes above, the reader may be able to see that the optimal projection may be obtained by evaluation of all of the constituent affine functions at each of the active cruxes, with a complexity-bound proportional to J^2. Due to space limitations, we do not include the details.

VI. Concluding Remarks

This paper considers a network disruption game, where the payoff derives from the reduced effectiveness of a player's controls caused by information flow delays. We find that the problem may be formulated with a min-plus convex cost and controlled min-plus linear dynamics. We also find that the solution may be obtained purely by idempotent linear operations. Complexity growth is the most significant issue. We find that the growth is severely limited by geometric considerations. We also find that complexity-growth attenuation is relatively straight-forward as well. More generally, computations are easily instantiated (and bear some similarity to those appearing in [8] for a different problem class), although space limitations prevent inclusion of examples here.

References

[1] M. Akian, "Densities of idempotent measures and large deviations", Trans. Amer. Math. Soc., 109 (1999), 79–111.
[2] G. Cohen, S. Gaubert, J.-P. Quadrat and I. Singer, "Max-Plus Convex Sets and Functions", preprint.
[3] P. Del Moral and M. Doisy, "Maslov idempotent probability calculus", Theory Prob. Appl., 43 (1999), 562–576.
[4] W.H. Fleming, H. Kaise and S.-J. Sheu, "Max-Plus Stochastic Control and Risk-Sensitivity", Applied Math. and Optim., 62 (2010), 81–144.
[5] S. Gaubert and W.M. McEneaney, "Min-max spaces and complexity reduction in min-max expansions", Applied Math. and Optim., (to appear).
[6] B. Heidergott, G.J. Olsder and J. van der Woude, *Max-Plus at Work: Modeling and Analysis of Synchronized Systems,* Princeton Univ. Press, 2006.
[7] V.N. Kolokoltsov and V.P. Maslov, *Idempotent Analysis and Its Applications,* Kluwer, 1997.
[8] W.M. McEneaney, "Idempotent Method for Deception Games and Complexity Attenuation", Proc. 2011 IFAC.
[9] W.M. McEneaney, "Idempotent Method for Deception Games", Proc. 2011 ACC, 4051–4056.
[10] W.M. McEneaney, "Complexity Reduction, Cornices and Pruning", Proc. of the Intl. Conf. on Tropical and Idempotent Mathematics, G.L. Litvinov and S.N. Sergeev (Eds.), Contemporary Math. 495, Amer. Math. Soc. (2009), 293–303.
[11] W.M. McEneaney, "Idempotent Method for Dynamic Games and Complexity Reduction in Min-Max Expansions", Proc. IEEE CDC 2009.
[12] A. Puhalskii, *Large Deviations and Idempotent Probability,* Chapman and Hall/CRC Press, 2001.
[13] A.M. Rubinov and I. Singer, "Topical and Sub-Topical Functions, Downward Sets and Abstract Convexity", Optimization, 50 (2001), 307–351.
[14] E.D. Sontag, "VC dimension of neural networks", Neural Networks and Machine Learning, Ed. C.M. Bishop, Springer (1998), 69–95.