

THE HEBREW UNIVERSITY OF JERUSALEM
FACULTY OF SCIENCE
EINSTEIN INSTITUTE OF MATHEMATICS

Timely Coordination in a Multi-Agent System

A research thesis submitted in partial fulfillment of the
requirements for the degree of Master of Science

Author:

Yannai A. Gonczarowski

Thesis Advisors:

Prof. Gil Kalai

Prof. Yoram Moses

Department of Electrical Engineering
Technion, Israel Institute of Technology

March 2012

To my parents, who taught me that knowledge is the one
thing that no one can ever take away from me,

to my brothers, who know everything,

and most of all, to Elee, the stable event of my life; the fixed
point of my life.

(This work would most probably never have been written were it not for her Kyoto adventure.)

Acknowledgements

I would like to thank my advisor, Prof. Gil Kalai, for his friendship and support throughout my studies. I would like to thank my second advisor, Prof. Yoram Moses, for his friendship, for introducing me to the fascinating subject of knowledge, and for his patience while I bombard him with emails during the small hours of the night.

In addition, I would like to thank my friends, my family, and last but not least, my other half Elee, for their support, their encouragement and their friendship, and simply for being there.

Abstract

In a distributed algorithm, multiple processes, or agents, work toward a common goal. More often than not, the actions of some agents are dependent on the previous execution (if not also on the outcome) of the actions of other agents. The resulting interdependencies between the timings of the actions of the various agents give rise to the study of methods for timely coordination of these actions.

In this work, we formulate and mathematically analyze a novel multi-agent coordination problem, which we call “Timely-Coordinated Response”, and in which the time difference between each pair of actions may be constrained by upper and/or lower bounds. This problem generalizes some classic coordination problems formulated and studied by Halpern and Moses, and some coordination problems recently formulated and studied by Ben-Zvi and Moses.

We optimally solve (i.e. provide an optimal protocol for solving) the timely-coordinated response problem in two ways: one using a generalization of the fixed-point approach of Halpern and Moses, and one using a generalization of the synchronous causality (“syncausality”) approach of Ben-Zvi and Moses. Furthermore, we constructively show the equivalence of the solutions yielded by both approaches, despite the vast conceptual differences between them. By combining both approaches, we derive strengthened versions of known results for some previously-defined special cases of this problem.

Our analysis is conducted under minimal assumptions: we work in a continuous-time model with possibly infinitely many agents. The general results we obtain for this model reduce to stronger results for discrete-time models with only finitely many agents. In order to distill the properties of such models that are significant to this reduction, we define several novel classes of naturally-occurring models, all generalizing discrete-time models with finitely many agents, which in a sense separate the different results. We investigate the timely-coordinated response problem in these models, and present both a more practical optimal solution for the problem, as well as a surprisingly simple condition for solvability thereof, for these models.

To conclude this work, we show how our results for the timely-coordinated response problem generalize the results known for previously-studied special cases of this problem, and present some open questions and further research directions.

Contents

1	Introduction	1
1.1	An Informal Example	1
1.2	Overview	3
2	Notation	6
3	A Discrete-Time Model	7
3.1	Context Parameters	7
3.2	Events and the Environment	8
3.3	States, Actions and Protocols	8
3.4	Runs	9
3.5	Notation	11
4	Timely-Coordinated Response	12
4.1	Coordinated Response Problems	12
4.2	Defining Timely-Coordinated Response	16
5	The Constraining Function	18
6	The Syncausality Approach	26
6.1	Background	26
6.1.1	Partial Orders on the Set of Agent-Time Pairs	26
6.1.2	Additional Notation	28
6.1.3	Previous Results	28
6.2	Adapting Some Machinery	34
6.3	Analyzing Timely-Coordinated Response	37

7	The Fixed-Point Approach	42
7.1	Background	42
7.1.1	Events, Knowledge and Common Knowledge	42
7.1.2	Fixed-Point Analysis	46
7.2	δ -Common Knowledge	48
7.3	Analyzing Timely-Coordinated Response	54
8	The Equivalence of Both Approaches	57
8.1	Background	58
8.2	A Constructive Proof	59
9	Results for Practical Models	66
9.1	Bounded-Syncausal-Path Contexts	66
9.2	Finite-Influence Protocols	75
9.3	Unbounded-Message-Delivery Contexts	81
10	Deriving Previous Results	83
10.1	General Ordered and Timed Responses	83
10.2	Common Knowledge and Variants	84
11	Discussion and Open Questions	87
11.1	A Qualitative Comparison of Approaches	87
11.2	On Generalizations	88
11.3	Open Questions and Further Directions	89
	Bibliography	91
	Appendices	
A	A Continuous-Time Model	95
A.1	Context Parameters	95
A.2	Timers	95
A.3	Agent States	96
A.4	Runs	97
A.5	Excluding Degeneracies	97
A.6	Modeling Discrete Time	99

List of Figures

1	Centipede	30
2	Broom	31
3	Centibroom	33
4	Path-traversing centipede	37
5	Path-traversing centipede with a broom suffix	71

Chapter 1

Introduction

In a distributed algorithm, multiple processes, or agents, work toward a common goal. More often than not, the actions of some agents are dependent on the previous execution (if not on the outcome) of the actions of other agents. This introduces interdependencies between the timing of the actions of the various agents.

1.1 An Informal Example

We begin with a simple example that illustrates how the coordination problem underlying this work arises as a natural, albeit nontrivial, continuation of previously studied problems.

Example 1.1. *Consider ACME, an IT company providing on-line storage services.* When ACME's on-line storage service is founded, its user base is relatively small, and one server fulfills all of the requirements of this service. To be on the safe side, though, ACME operates a backup server. Being fairly confident in the stability of its main server, ACME does not impose any freshness constraints on the backup server, other than that the backup server must never be ahead of the main server, in order to avoid even the slightest potential of a data change being reflected solely in the backup server. Thus, we may concisely capture the only timing constraint ACME imposes on its servers: If a user changes her data, then eventually the main server*

* Some readers may be acquainted with other products provided by ACME, such as high-tech warfare gear used by Wile E. Coyote in his endless quest to capture the Road Runner, or be acquainted with the ACME detective agency on the hunt for V.I.L.E. ringleader and former ACME agent Carmen Sandiego. As both these venues have become less lucrative in recent years, ACME decided to follow the trend and go into IT.

reflects this change, and eventually, at some later time, the backup server reflects it as well. A generalization of the problem underlying such a scenario was studied by Lamport[17] in an asynchronous model. Recently, Ben Zvi and Moses[6, 5, 7] extended this study to synchronous models as well, dubbing the generalized problem “ordered response”.

We return to ACME’s story. After a while, the user base of the company’s storage service grows and moreover, many users use it more heavily than before, as they have grown both accustomed to it, and confident of its abilities and stability. Eventually, a single-server solution becomes inadequate for this service, and ACME turns its backup server into a second live server. Optimally, ACME would like to impose the following timing constraint on its servers: If a user changes her data, then eventually both servers reflect this change, and they do so simultaneously. A generalization of the problem underlying such a scenario has been extensively studied[21, 10] under the name “firing squad”. In particular, it was studied both by Ben-Zvi and Moses[6, 5, 7], who dub it “simultaneous response”, and by Halpern and Moses[14, 11], who call it “perfect coordination”. Unfortunately for ACME, though, it is shown in [14] that this problem is unsolvable under realistic conditions.

Having read [14], ACME decides to go for what its engineers perceive as “the next best thing”, replacing the requirement for simultaneous reflection of a change in both servers to “almost simultaneous” reflection. Formally, they demand that if any server reflects a change at any time, then the other server reflects this change no later than 100 milliseconds thereafter. The problem underlying such a scenario no longer falls within the scope of the study of Ben-Zvi and Moses[6, 5, 7, 8], although a generalization thereof was studied by Halpern and Moses[14, 11] under the name “ ϵ -coordination”.

Naturally, as long as this “near simultaneity” constraint (or, in the preceding scenarios, the relevant timing constraint introduced there) is met, ACME wish for both servers to reflect each user action as close to the time of its occurrence as possible.*

Shortly after the switch to two live servers, an email reaches ACME’s headquarters. ACME’s on-line gaming subsidiary, which uses ACME’s on-line storage infrastructure to store the state of their on-line multi-player games, complains that the slow response time of ACME’s servers, coupled with a 100-millisecond lag between

* We later phrase both the worst-case response time in each of the above scenarios, as well as a necessary and sufficient condition for the sheer solvability of this problem, as functions of the topology of ACME’s network, and of the worst-case communication lag times in it.

these servers, renders its multi-player games unplayable. ACME engineers convene for an emergency meeting, and propose the following plan: as most of ACME's gaming customers are located in the same vicinity (ACME's gaming platform is very popular in Israel), from which traffic to ACME's server #2 is particularly fast (server #2 resides in Israel, while server #1 resides in the U.S.A.), they can simply route all the gaming traffic to server #2, effectively eliminating the 100-millisecond lag the gaming subsidiary is complaining about. Unfortunately, this does not solve the other cause of this complaint: the slow response time of ACME's servers. (The performance price of ACME's new algorithm, which coordinates a maximum freshness lag of 100 milliseconds between the servers, is a slower response time of both servers than the one achieved in the old single-live-server algorithm.) One of the engineers raises the following question: perhaps they can achieve a faster response time for server #2 if the timing constraint is revised to be asymmetric: server #2 must update no later than 100 milliseconds after server #1, however server #1 may update as late as 300 milliseconds after server #2. The problem underlying such a scenario falls out of the scope both of the studies of Ben-Zvi and Moses[6, 5, 7, 8], and the studies of Halpern and Moses[14] and their extension by Fagin et al.[11, Section 11.6].

The multi-agent coordination problem that we present and analyze in this thesis generalizes, among others, the problems arising in the above example. In particular, the problem that we present generalizes the last problem arising from this example, by allowing to arbitrarily bound the time difference between each pair of actions both from above and from below. This generalizes the study of Halpern and Moses[14] (and of Fagin et al.[11, Section 11.6]) by allowing different bounds to be specified for different pairs of actions, and generalizes the study of Ben-Zvi and Moses[6, 5, 7, 8] by allowing the specification of an upper and a lower bound that do not coincide, on the time difference between a pair of actions.

1.2 Overview

In this work, we present and mathematically analyze a novel multi-agent coordination problem, which we call “timely-coordinated response”. In this coordination problem, which we define and analyze in a synchronous model, a set of agents are to perform local actions, and the time difference between each pair of actions may be constrained by an upper and/or a lower bound (or neither), which are param-

ters given as part of the problem description. Following the studies of Halpern and Moses[14], Fagin et al.[11], and Ben-Zvi and Moses[6, 5, 7, 8], which we generalize, most of this work revolves around the interaction between time and coordination.

After presenting the timely-coordinated response problem in Chapter 4, we perform, in Chapter 5, a graph-theoretical analysis of the set of constraining parameters (upper and lower bounds on the time difference, for each pair of actions) that define this problem. This analysis leads to a definition of a canonical representative for each class of constraint-sets that define the same problem, and to a characterisation for solvability of the timely-coordinated response problem under what may be regarded as ideal conditions.

In the following two chapters, we optimally solve the timely-coordinated response problem in two ways, each generalizing one of the approaches previously used to analyze some special cases thereof: In Chapter 6, we survey, and then generalize, the “syncausality” approach of Ben-Zvi and Moses, which may be viewed as more of a concrete “nuts and bolts” approach. In this chapter, which is combinatorial in character, we study the timely-coordinated response problem in the possible presence of guarantees on message delivery times between agents. (In Chapter 9, we present a result showing the impossibility of timely coordination using mutual constraints in the absence of such bounds.) In Chapter 7, we survey, and then generalize, the “fixed-point” approach of Halpern and Moses, which was later studied by Fagin et al. as well. This approach, whose origins are traceable to temporal logic, may be conversely viewed as more of an abstract “higher level” approach. The main result presented in each of these two chapters is a description of an optimal protocol/algorithm for solving the timely-coordinated response problem. Following these analyses, we constructively show, in Chapter 8, that, despite the significant conceptual and technical differences between these two approaches, they both yield equivalent solutions for the timely-coordinated response problem.

The above-surveyed analysis is conducted under minimal assumptions: it applies to a continuous-time model, which may contain infinitely many agents. The general results obtained using this analysis reduce to stronger results when the model in question is a discrete-time one, and contains only finitely many agents. In Chapter 9, we define several novel classes of naturally-occurring models, all of which generalize discrete-time models that contain finitely many agents. These classes of models, in a sense, separate the specialized discrete-time finite-agent results from the generic continuous-time infinite-agent results. We investigate the

timely-coordinated response problem in these models, and derive, for these models, both a more practical description of the optimal solution for this problem, as well as a surprisingly simple condition for solvability thereof, in terms of the available network communication channels and the worst-case delivery times therein. We conclude this chapter by combining both approaches to derive a strengthened version of a known impossibility result for some previously-defined special cases of the timely-coordinated response problem.

Following the above analysis, we show, in Chapter 10, how the results obtained in the previous chapters reduce to generalizations of the known best results for previously-studied special cases of the timely-coordinated response problem.

Finally, in Chapter 11, we qualitatively discuss some of our results, and present some open questions and some future research directions in which the results of this work may prove to be useful.

The main contributions of this work are:

1. Identifying, defining and analyzing the timely-coordinated response problem.
2. Applying both above-described approaches to analyze this problem, thereby unifying, generalizing and strengthening results previously achieved using these approaches.
3. Deriving generic results for continuous-time models, as well as specialized results for discrete-time models, and defining “intermediate” model classes which, in a sense, separate the continuous-time results from the discrete-time results.

Underlying this work are three different currents. While these are interconnected, each of these may stand alone in its own right, and may be of interest to a different audience:

- Our graph-theoretic analysis from Chapter 5 may be of most interest to combinatorists.
- Our generalization of the fixed-point approach and its results may be of most interest to logicians and game theorists.
- Our generalization of the syncausality approach and its results may be of most interest to computer scientists and engineers.

Chapter 2

Notation

Throughout this work, we use the following notation:

- $\forall n \in \mathbb{N} : [n] \triangleq \{1, \dots, n\}$.
- We denote the non-negative reals by $\mathbb{R}_{\geq 0} \triangleq \{t \in \mathbb{R} \mid t \geq 0\}$.
- Given a set I , we denote the set of ordered pairs of distinct elements of I by

$$I^2 \triangleq \{(i, j) \in I^2 \mid i \neq j\}.$$

- Given a set A , $n \in \mathbb{N}$ and an n -tuple $\bar{a} = (a_m)_{m=1}^n \in A^n$, we denote the n -tuple containing the elements of \bar{a} in reverse order by $\bar{a}^{rev} \triangleq (a_{n-m+1})_{m=1}^n$.
- Given a directed graph $G = (V, E)$, we denote the set of paths in G by

$$\mathcal{P}(G) \triangleq \{\bar{p} = (p_m)_{m=1}^n \in V^n \mid n \in \mathbb{N} \ \& \ \forall m \in [n-1] : (p_m, p_{m+1}) \in E\}.$$

- Given a weighted directed graph $G = (V, E, w)$, we denote the length of a path $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G)$ by

$$L_G(\bar{p}) \triangleq \sum_{m=1}^{n-1} w(p_m, p_{m+1}).$$

Furthermore, we denote the distance function between vertices of G by

$$\begin{aligned} \delta_G : V^2 &\rightarrow [-\infty, \infty] \\ (i, j) &\mapsto \inf\{L_G(\bar{p}) \mid \bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G) \ \& \ p_1 = i \ \& \ p_n = j\}. \end{aligned}$$

Chapter 3

A Discrete-Time Model

We model a set of agents that communicate with each other solely via message passing. Each agent follows a predetermined protocol, which is common knowledge to all agents. In the following chapters, we concern ourselves with the task of devising such protocols with the goal of analyzing a coordination problem that we define in Chapter 4.

To avoid over-burdening the reader with cumbersome details, the model presented in this chapter is a discrete-time model, conceptually based on [11], which may be assumed while reading this work. It should be noted, though, that the results presented throughout this work hold verbatim also for a more intricate continuous-time model, which we present in Appendix A. In Chapter 9, we consider several natural properties of practical continuous-time models, some of which always hold for the model presented in this chapter, and prove results for models with these properties.

3.1 Context Parameters

Intuitively, a context describes the environment in which the agents operate. In this discrete-time model, we formally denote a context by a tuple $\gamma \triangleq (\mathcal{G}_\gamma = (\mathbb{I}_\gamma, N_\gamma, b_\gamma), (S_i)_{i \in \mathbb{I}_\gamma}, \tilde{E}_\gamma, (i_{\tilde{e}})_{\tilde{e} \in \tilde{E}_\gamma})$, where:

1. \mathcal{G}_γ is a weighted directed graph with positive, integral or infinite, weights. The vertices \mathbb{I}_γ of \mathcal{G}_γ model the agents. We say that an agent $j \in \mathbb{I}_\gamma$ neighbours an agent $i \in \mathbb{I}_\gamma \setminus \{j\}$ if $(i, j) \in N_\gamma$. If this is a case, then i may, as will be defined in greater precision below, send messages to j , which are guaranteed to arrive

no later than $b_\gamma(i, j)$ after being sent.

2. For each agent $i \in \mathbb{I}_\gamma$, S_i is a set of legal states for i . We assume that S_i is of large enough cardinality to accommodate all our needs.
3. \tilde{E}_γ is a set of “possible external inputs”. We think of external inputs as non-deterministic events, the occurrence of which may not be anticipated in advance by any agent.
4. Each external input $\tilde{e} \in \tilde{E}_\gamma$ is associated with a single agent $i_{\tilde{e}} \in \mathbb{I}_\gamma$, which observes this input when it occurs.

Additionally, we define the set of times as $\mathbb{T} \triangleq \mathbb{N} \cup \{0\}$. As noted above, in Appendix A we give an alternative model description, in which time is continuous.

3.2 Events and the Environment

At each possible time $t \in \mathbb{T}$, zero or more events may take place. Intuitively, an event is an occurrence that is observed by a single agent. An event may be of one of the following types:

1. An external input event $\tilde{e} \in \tilde{E}_\gamma$. (Observed by $i_{\tilde{e}}$.)
2. A message delivery event $(m, t', (i, j))$ of a message m , sent by i at time t' , to j , s.t. $(i, j) \in N_\gamma$. (Observed by j .)

We define the “state of the environment” at any given time as the set of events that take place at that time. We denote the set of all possible states of the environment by $S_e \triangleq 2^{\tilde{E}_\gamma} \times 2^{\mathcal{M} \times \mathbb{T} \times N_\gamma}$, where \mathcal{M} is a set of all possible messages.

3.3 States, Actions and Protocols

The problem that we define in the next chapter deals with the coordination of the responses of different agents to an external input. At any time $t \in \mathbb{T}$, each agent $i \in \mathbb{I}_\gamma$ performs the following, in a manner that is based on its state, as well as on any events observed by it at t :

1. Sets a new state for itself (which may be identical to its old state).

2. Sends any number of messages, each with possibly different content, and to a possibly different neighbouring agent $j \in \mathbb{I}_\gamma$ (i.e. $j \in \mathbb{I}_\gamma$ s.t. $(i, j) \in N_\gamma$).
3. Possibly “responds”. This is the action that we aim to coordinate.*

We thus define the set of possible actions that may be taken by i at t as $A_i \triangleq S_i \times 2^{\mathcal{M} \times \{j \in I \mid (i, j) \in N_\gamma\}} \times \{\text{false}, \text{true}\}$. (Each element in A_i consists of a new state for i at t , a set of messages to be sent by i at t , and a boolean value that indicates whether i responds at t .)

A “local protocol” for an agent $i \in \mathbb{I}_\gamma$ consists of a set of possible initial states for i , together with an “action” function, receiving as input a state of i just before a certain time $t \in \mathbb{T}$ and any events observed by i at t , and outputting the actions to be performed by i at t .[§] Formally, a local protocol for i is a pair (\tilde{S}_i, P_i) , s.t. $\tilde{S}_i \subseteq S_i$ and $P_i : S_i \times 2^{\{\tilde{e} \mid i_{\tilde{e}} = i\}} \times 2^{\mathcal{M} \times I} \rightarrow A_i$. In certain cases, we may wish to allow the actions of i at t to also depend on t as well. In such cases, we say that the model is a “shared-clock model”, and the actions function takes the form $P_i : S_i \times 2^{\{\tilde{e} \mid i_{\tilde{e}} = i\}} \times 2^{\mathcal{M} \times I} \times \mathbb{T} \rightarrow A_i$.

A “joint protocol” (“protocol”, for short) is a collection of local protocols, one for each $i \in \mathbb{I}_\gamma$. We denote the set of all protocols of γ by \mathbb{P}_γ .

An important set of protocols is the set of “full-information” protocols[20], in which the state of each agent $i \in \mathbb{I}_\gamma$ at any time $t \in \mathbb{T}$ uniquely determines the full details of every event observed by i up until, and including, t . Furthermore, at every $t \in \mathbb{T}$, i sends to every neighbouring agent a message including the full current state of i at t , and in a shared-clock model — also the current time.

3.4 Runs

For the duration of this section, fix a protocol $P = ((\tilde{S}_i, P_i))_{i \in I} \in \mathbb{P}_\gamma$. A “run” of P in γ is, intuitively, a “possible infinite history” of P executed in γ , in which the behaviour of all agents is governed by P . Formally, a run of P is a function $r : \mathbb{T} \rightarrow S_e \times \prod_{i \in I} S_i$ assigning, for each time, a state for each agent, and a state for the environment, while satisfying the following properties:[¶]

* For simplicity, we define only one type of response per agent. Our models, and our results in this work, may be readily generalized to allow a set of possible responses for each agent.

[§] We restrict ourselves to deterministic protocols solely for ease of exposition.

[¶] While a run is customarily defined in an inductive fashion, we define it here without induction in order to minimize the differences between the definitions of our discrete- and continuous-time models.

- Agent state consistency with local protocol: Let $i \in I$ and $t \in \mathbb{T}$, then $r_i(t)$ must equal the first part of the output of P_i when applied to $r_i(t-1)$ (or to some $\tilde{s}_i \in \tilde{S}_i$, if $t = 0$) and to the events observed by i at t (which depend only on $r_e(t)$). (The other parts of this output of P_i determine the actions of i at t .)
- Environment state properties:
 1. Each external input $\tilde{e} \in \tilde{E}_\gamma$ may occur no more than once during a run.
 2. Bounded message delivery: If a message is sent at any $t \in \mathbb{T}$ by $i \in \mathbb{I}_\gamma$ to $j \in \mathbb{I}_\gamma$ (where $(i, j) \in N_\gamma$), then it must be delivered exactly once, at some time $t' \in \mathbb{T}$, s.t. $t < t' \leq t + b_\gamma(i, j)$. If $t' < t + b_\gamma(i, j)$, then we say that this message is delivered *early*. Only messages that are sent during a run may be delivered during it.

We call an event “non-deterministic” (ND for short) if it is either an external input event, or an early delivery. (Intuitively, we may think of ND events as events that cannot be foreseen by any agent before they occur, and thus occur, in a sense, at the whim of the environment.) It should be noted that this definition is context-dependent, as it depends on b_γ .

We denote the set of all possible runs of P in γ by $R_\gamma(P)$. We denote the set of all runs of all protocols by in γ by $\mathcal{R}_\gamma \triangleq \cup_{P \in \mathbb{P}_\gamma} R_\gamma(P)$.

It should be noted that two full-information protocols (as defined in the previous section) $P, P' \in \mathbb{P}_\gamma$ may only differ in their response logic, and therefore there is a natural isomorphism between $R_\gamma(P)$ and $R_\gamma(P')$, which preserves the set of ND events along with their occurrence times.

As was essentially shown in [20], given a protocol $P \in \mathbb{P}_\gamma$, there exists a full-information protocol $P' \in \mathbb{P}_\gamma$, s.t. there is a natural monomorphism from $R_\gamma(P)$ into $R_\gamma(P')$, which preserves both the set of ND events (and their occurrence times), and all responses (and response times). This ability of a full-information protocol to simulate any other protocol implies, for our purposes, that if there exists some protocol that solves a certain coordination problem, then there also exists a full-information protocol that solves this problem. This justifies restricting to full-information protocols when analyzing solvability, which will sometimes prove convenient. (See, e.g. Definition 4.7 and Corollary 9.7.)

3.5 Notation

Given a context $\gamma = (\mathcal{G}_\gamma = (\mathbb{I}_\gamma, N_\gamma, b_\gamma), (S_i)_{i \in \mathbb{I}_\gamma}, \tilde{E}_\gamma, (i_{\tilde{e}})_{\tilde{e} \in \tilde{E}_\gamma})$, we introduce the following notation:

- Given a run $r \in \mathcal{R}_\gamma$, we denote the set of all events in r by $E(r)$. As we wish to regard each event as unique, and as we defined events above not to contain their occurrence time, we technically define $E(r)$ as a set of event-time pairs (e, t_e) , where each event is paired with its occurrence time.
- Given a run $r \in \mathcal{R}_\gamma$, we denote the set of all the ND events in r by $ND_\gamma(r) \subseteq E(r)$.
- As mentioned above, we analyze the coordination of the responses of different agents to a given external input. Thus, given a protocol $P \in \mathbb{P}_\gamma$ and an external input $\tilde{e} \in \tilde{E}_\gamma$, we define the set of “ \tilde{e} -triggered” runs of P as

$$R_\gamma^{\tilde{e}}(P) \triangleq \{r \in R_\gamma(P) \mid \tilde{e} \in ND_\gamma(r)\}.$$

where by slight abuse of notation, we use $\tilde{e} \in ND_\gamma(r)$ as a shorthand for $\exists t \in \mathbb{T} : (\tilde{e}, t) \in ND_\gamma(r)$. We follow this convention occasionally, when no confusion can arise.

Chapter 4

Timely-Coordinated Response

In this chapter, we define the main coordinated response problem underlying this work, which we call “timely-coordinated response”, and analyze some of its basic properties.

4.1 Coordinated Response Problems

Before defining the timely-coordinated response problem, we first define a broader, much simpler, coordinated response problem. (In a sense, it is the simplest coordinated response problem.) This simpler problem will serve as a building block for the timely-coordinated response problem. In addition, the short discussion of this simpler problem will provide an introduction to a novel theory of coordinated response problems,* which we attempt to formalize throughout this work. During this discussion, we introduce some definitions and concepts, which we later reuse while discussing various coordinated response problems throughout this work. While philosophically a problem may be thought of as a specification, or as a collection of constraints, we formally associate a coordinated response problem with the set of protocols solving it, effectively treating two problems that share the same set of solutions as identical.

Definition 4.1 (Eventual Response). *Given a context γ , an external input $\tilde{e} \in \tilde{E}_\gamma$ and a set of agents $I \subseteq \mathbb{I}_\gamma$, we define the “eventual response” problem $ER_\gamma\langle\tilde{e}, I\rangle \subseteq$*

* While Ben-Zvi and Moses[6, 5, 7, 8] define quite a few problems to which they refer as “response problems” (all of which we survey in Chapters 6 and 10 and from which we draw motivation), they do not give a formal definition for a response problem, or for a coordinated response problem.

\mathbb{P}_γ as the set of all protocols P satisfying:

- In each $r \in R_\gamma^\tilde{e}(P)$, each $i \in I$ responds exactly once. In this case, we denote, for each $i \in I$, the response time of i in r by $t_r(i)$. (Hence, t_r is a function from I to \mathbb{T} .)
- In each $r \in R_\gamma(P) \setminus R_\gamma^\tilde{e}(P)$, neither of the agents in I responds. In this case, we define $t_r \equiv \infty$.

Thus, for every $r \in R_\gamma(P)$, we have defined a function $t_r : I \rightarrow \mathbb{T} \cup \{\infty\}$.

Remark 4.2. Let $P \in ER_\gamma\langle\tilde{e}, I\rangle$ and let $r \in R_\gamma^\tilde{e}(P)$. While each $i \in I$ responds in r , none of them do so before \tilde{e} occurs.

Intuitively, Remark 4.2 holds because, as we noted when describing our model(s), the non-deterministic nature of \tilde{e} implies that no agent may possibly infer that a run is triggered before \tilde{e} occurs in that run. (Indeed, as far as any agent is concerned, as long as \tilde{e} has not occurred yet, it may be the case that \tilde{e} will never occur during this run.*) Formally, this may be readily proven using machinery that we have not introduced yet. (See, e.g. the beginning of the proof of Theorem 6.34 for a proof of a stronger statement using the tools of Chapter 6, and Corollary 7.30 for a conceptually-similar proof using the tools of Chapter 7.)

Example 4.3. In the popular TV talent show “Got Talent”, a panel of judges (that we denote by I) judge various amateur performances, with themes ranging from music, through magic, to some very obscure themes that are better left undescribed. Once a contestant starts performing on stage, each judge may press an “X” button to signal her desire for this performance to end. (We denote this action as a response by the judge.) The performance continues until it has run its course, or until all judges have pressed their respective “X” buttons.

As the “mean” judges sometimes implicitly compete between themselves regarding who presses his “X” button first, let us consider a hypothetical “enhancement” to the show, in which a bucket of water is poured over the head of any judge who presses his “X” button before a performance starts.

Consider a hypothetical repetitive (and thus, potentially never-ending) performance in this TV show. (This indeed sometimes seems to be the case, especially

* In the continuous-time model presented in Appendix A, this stems from the “no foresight” property.

for very bad performances, which are abundant in this show.) In order to guarantee that this performance does not continue forever, the show producers must make sure that each judge presses her “X” button at some time during the performance (but not before the performance starts, as she wishes to remain dry). The order of the “X” presses of the different judges is insignificant, as long as it is guaranteed that each judge eventually presses her “X” button. As the beginning of the performance depends on the rambling of the comedian serving as show host (which may possibly also never end), it is impossible for the judges to predict before it actually occurs, so we may regard it as a non-deterministic external input, which we denote as \tilde{e} .

Remark 4.4. In Definition 4.1 and hereafter, we assume, for simplicity, that each agent is associated with no more than one response, as we did when defining our model. (While it may seem a bit silly to assume otherwise in the eventual response problem, it may make sense to do so for more intricate coordinated response problems defined below.) Nonetheless, all the results we derive throughout this work regarding any response problem are easily adaptable to the case in which more than one response type is available per agent, and in which I is a set of agent-response pairs, rather than merely agents. (In this way, some agents may be associated with more than one of the responses being coordinated.) Indeed, all our results apply verbatim to this generalized case as well, if we allow ourselves, for each agent-response pair $i = (\tilde{i}, a) \in I$, to slightly abuse notation by writing i to refer to \tilde{i} as well.

While studying a coordinated response problem, we are usually interested in two questions:

- Solvability: Under which conditions is it solvable? and if so,
- Optimality: What is the “fastest” way to solve it?

Before answering these questions regarding the eventual response problem, we first define them (and define coordinated response problems) precisely.

Definition 4.5 (Coordinated Response). Let γ be a context, let $\tilde{e} \in \tilde{E}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. We call a problem a “coordinated response” problem to \tilde{e} by I , if the set of protocols solving it is a subset of (the set of protocols solving) $ER_\gamma\langle\tilde{e}, I\rangle$. As before, we formally identify a coordinated response problem with the set of protocols solving it.

We usually define coordinated response problems by restrictions on the response times t_r in all triggered runs r (e.g. “ I must respond together in each triggered run”,

“ I must respond in a given order in each triggered run”, etc.). In order to ease the reading of the following two definitions, one may consider, as an example, the case in which $CP \triangleq ER_\gamma\langle\tilde{e}, I\rangle$ for some $\tilde{e} \in \tilde{E}_\gamma$ and $I \subseteq \mathbb{I}_\gamma$.

Definition 4.6 (Solvability). *Let γ be a context and let $CP \subseteq \mathbb{P}_\gamma$ be (the set of protocols solving) a coordinated response problem. We say that CP is “solvable” if $CP \neq \emptyset$. Otherwise, we say that it is “unsolvable”.*

Definition 4.7 (Optimal Response Logic). *Let γ be a context and let $CP \subseteq \mathbb{P}_\gamma$ be a coordinated response problem. Assume that CP is solvable and let $P \in CP$. Assume w.l.o.g. that P is a full-information protocol. Let P' be the protocol obtained from P by modifying the response logic of each $i \in I$ in some way. As P and P' only differ by their response logic, there is a natural isomorphism between $R_\gamma(P)$ and $R_\gamma(P')$ that preserves the set of ND events.*

We say that the response logic of P' is an “optimal response logic” for solving CP if the following are satisfied:

1. $P' \in CP$.
2. *If $r \in R_\gamma(P)$ and $r' \in R_\gamma(P')$ are two runs of the respective protocols matched under the above isomorphism, then $t_{r'} \leq t_r$.*

We may now answer the above questions of solvability and optimality, with regard to the eventual response problem.

Remark 4.8. *Let γ be a context, let $\tilde{e} \in \tilde{E}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. The following may be easily verified:*

- $ER_\gamma\langle\tilde{e}, I\rangle$ is solvable iff for every $i \in I$ there exists $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G)$ s.t. $p_1 = i_{\tilde{e}}$ and $p_n = i$.
- An optimal response logic for solving $ER_\gamma\langle\tilde{e}, I\rangle$ is, for every $i \in I$, “respond as soon as i receives information guaranteeing that \tilde{e} occurred.”* (This seemingly-vague condition has a very precise meaning in a full-information protocol, as each message sent in such a protocol uniquely determines a list of events that are guaranteed to have occurred.)

* In some runs of certain contexts with infinitely many agents under the continuous-time model presented in Appendix A, the set of times at which i has information guaranteeing that \tilde{e} has occurred does not attain its infimum value. It is straightforward to show that for such pathological cases, no optimal response logic exists. Similar observations hold for all other optimal response logics presented in this work as well.

It should be noted that for many coordinated response problems, it is fairly straightforward to deduce a solvability criterion (phrased as a requirement on the topology and weights of \mathcal{G}_γ) from an optimal response logic. In the above example, the optimal response logic demands, for i to respond in each triggered run, that in each such run i receives information guaranteeing that \tilde{e} occurred. Thus, a necessary and sufficient condition for solvability is that in each triggered run, each agent is guaranteed to receive information to this effect. This is exactly equivalent to the solvability criterion above. For this reason, for many coordinated response problems, we will be primarily interested in an optimal response logic.

Due to the first part of Remark 4.8, we assume hereafter, whenever discussing a coordinated response problem of a set of agents $I \subseteq \mathbb{I}_\gamma$ to an external input $\tilde{e} \in \tilde{E}_\gamma$, that there exist paths in \mathcal{G}_γ from $i_{\tilde{e}}$ to every agent $i \in I$.

4.2 Defining Timely-Coordinated Response

We define the timely-coordinated response problem as a response problem in which maximum and minimum values for the differences between the response times of each pair of agents are provided. In a discrete-time model, such constraints may be defined by any integral, or infinite, value. In a continuous-time model, such constraints may be defined by any real, or infinite, value. In order to formally present the timely-coordinated response problem, we first define the set of such possible constraining values.

Definition 4.9. We define $\Delta = (\mathbb{T} - \mathbb{T}) \cup \{-\infty, \infty\}$.

Remark 4.10. If $\mathbb{T} = \mathbb{N} \cup \{0\}$, then $\Delta = \mathbb{Z} \cup \{-\infty, \infty\}$. If $\mathbb{T} = \mathbb{R}_{\geq 0}$, then $\Delta = [-\infty, \infty]$.

We now turn to define the constraints imposed on response times in triggered runs in the timely-coordinated response problem. Recall that given a set I , we denote the set of ordered pairs of distinct elements of I by I^2 .

Definition 4.11 (Implementation). We call a pair (I, δ) an “implementation-spec”, if I is a set and if δ is a function $\delta : I^2 \rightarrow \Delta$. Given an implementation-spec (I, δ) , we call a function $t : I \rightarrow \mathbb{T}$ an “implementation” of δ , if $t(j) \leq t(i) + \delta(i, j)$ for every $(i, j) \in I^2$. We denote the set of all implementations of δ by $T(\delta)$. If $T(\delta) \neq \emptyset$, we say that δ is “implementable”. Otherwise, we say that it is “unimplementable”.

Remark 4.12. Let (I, δ) be an implementation-spec. By Definition 4.11:

- Obviously, δ is unimplementable unless $\delta > -\infty$. Nonetheless, we still allow δ to take the value of $-\infty$ for some or all agent pairs, for technical reasons that may become apparent when we define a canonical form for δ in the next chapter.
- Every $t \in T(\delta)$ satisfies $-\delta(j, i) \leq t(j) - t(i) \leq \delta(i, j)$ for every $(i, j) \in I^2$.
- Let $t : I \rightarrow \mathbb{T}$. If $t \in T(\delta)$, then $t + c \in T(\delta)$ as well, for every $c \in \mathbb{T}$. We say that two implementations of δ are “similar” if they differ by a translation.
- T is monotone: Let $\delta' : I^2 \rightarrow \Delta$. If $\delta \leq \delta'$, then $T(\delta) \subseteq T(\delta')$.

At last, we are ready to define the timely-coordinated response problem.

Definition 4.13 (Timely-Coordinated Response). We call a quadruplet $(\gamma, \tilde{e}, I, \delta)$ a “TCR-spec”, if γ is a context, $\tilde{e} \in \tilde{E}_\gamma$ is an external input, and (I, δ) is an implementation-spec s.t. $I \subseteq \mathbb{I}_\gamma$. Given a TCR-spec $(\gamma, \tilde{e}, I, \delta)$, we define the “timely-coordinated response” problem $TCR_\gamma\langle\tilde{e}, I, \delta\rangle \subseteq ER_\gamma\langle\tilde{e}, I\rangle$ as the set of all eventual-response protocols P for which $t_r \in T(\delta)$ for every triggered run $r \in R_\gamma^\tilde{e}(P)$.

Remark 4.14. Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec. By Definition 4.13:

- $\forall J \subseteq I : TCR_\gamma\langle\tilde{e}, I, \delta\rangle \subseteq TCR_\gamma\langle\tilde{e}, J, \delta|_{J^2}\rangle$.
- Let $\delta' : I^2 \rightarrow \Delta$. If $T(\delta) \subseteq T(\delta')$, then $TCR_\gamma\langle\tilde{e}, I, \delta\rangle \subseteq TCR_\gamma\langle\tilde{e}, I, \delta'\rangle$.

Example 4.15. Returning to the “Got Talent” show from Example 4.3. Assume that the panel of judges consists of two judges: Alice and Bob. The producers wish to create, among the viewers, the general impression that Bob is a “meaner” judge than Alice, but only by a subtle difference. One way to achieve this may be to ensure that Bob never responds more than 5 seconds after Alice, and that Alice never responds more than 40 seconds after Bob. Coordinating this may not be very simple if, for example, the judges are seated in a way that prevents each judge from knowing when the other judge presses the “X” button. (Thus, for example, Alice may have to rely on information regarding Bob’s taste, such “Bob always presses his “X” button no more than 10 seconds after someone falls on stage.”) This is a simple instance of the timely-coordinated response problem.

The rest of this work, as noted above, is dedicated to the analysis of the timely-coordinated response problem.

Chapter 5

The Constraining Function

Before we turn to analyze the coordination required in order to solve the timely-coordinated response problem, we first note that for an unimplementable δ , the situation is hopeless to begin with, as the timely-coordinated response problem is unsolvable regardless of the context γ in which it is defined. In this chapter, we embark on a graph-theoretic discussion with the aim of phrasing a necessary and sufficient condition for implementability of a constraining function δ . First, though, we make the above comment regarding “hopelessness” precise:

Claim 5.1. *Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec.*

1. *$TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is unsolvable if δ is unimplementable.*
2. *If $TCR_\gamma\langle\tilde{e}, I, 0\rangle^*$ is solvable, then the converse holds as well, i.e. $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable if δ is implementable. Furthermore, for any implementation \tilde{t} of δ , there exists a solving protocol $P \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle$, for which the map from agents to response times in every one of its triggered runs is similar to \tilde{t} .*

Remark 5.2. *Regarding Claim 5.1:*

- *Solvability of $TCR_\gamma\langle\tilde{e}, I, 0\rangle$ is equivalent to the ability to coordinate a simultaneous response of all agents in I in every \tilde{e} -triggered run. This classic problem, known as the “Firing Squad” problem[21, 10] will repeatedly appear in this work. (See, e.g. Theorem 6.12, Corollary 9.7, Theorem 9.10 and Theorem 9.12.)*

* We use 0 here and hereafter as a shortcut for the constant zero constraining function, i.e. $\delta_0 : I \rightarrow \Delta$ s.t. $\delta_0 \equiv 0$.

- *Corollary 9.7 shows that under certain conditions, the 0 function in the second part of Claim 5.1 may be replaced with a variety of other functions.*

Proof of Claim 5.1. For the first part, assume that $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable. Thus, there exists $P \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle$. Let $r \in R_\gamma^\tilde{e}(P)$. By Definition 4.13, t_r is an implementation of δ , and hence δ is implementable.

For the second part, assume that δ is implementable and let $\tilde{t} \in T(\delta)$. Let $P_0 \in TCR_\gamma\langle\tilde{e}, I, 0\rangle$. Let P be the protocol obtained from P_0 by modifying the response logic of each $i \in I$ to “respond $\tilde{t}(i)$ time units after the time i would have responded in P_0 ”. (This may require adding some auxiliary variables, which consume only a finite amount of memory, to the state of i .^{*} For the continuous-time model presented in Appendix A, this response logic may be implemented using timers.) We complete the proof by showing that $P \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle$. As P_0 and P only differ by their response logic, there is a natural isomorphism between $R_\gamma(P_0)$ and $R_\gamma(P)$, which preserves the set of ND events. Let $r \in R_\gamma(P)$, and denote by r_0 the run of P_0 matched to r under this isomorphism. If $r \in R_\gamma^\tilde{e}(P)$, then $r_0 \in R_\gamma^\tilde{e}(P_0)$, and we have $t_r = t_{r_0} + \tilde{t} < \infty$. Note that as $P_0 \in TCR_\gamma\langle\tilde{e}, I, 0\rangle$, we obtain that t_{r_0} is a constant function — denote its value by t_{r_0} . By Remark 4.12, $t_r = t_{r_0} + \tilde{t}$ is an implementation of δ (which is, by definition, similar to \tilde{t}). If $r \notin R_\gamma^\tilde{e}(P)$, then $r_0 \notin R_\gamma^\tilde{e}(P_0)$. In this case, we have $t_{r_0} \equiv \infty$, and therefore $t_r \equiv \infty$ as well. \square

By the second part of Claim 5.1, the study of the implementability of a constraining function δ may also be thought of as the study of solvability of $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ in contexts γ in which $TCR_\gamma\langle\tilde{e}, I, 0\rangle$ is solvable. (By both parts of that claim, contexts in which $TCR_\gamma\langle\tilde{e}, I, 0\rangle$ is solvable may be thought of as ideal for solvability, in the sense that instances of the timely-coordinated response problem that are unsolvable therein are unsolvable in any context.[§])

As a first step toward analyzing the implementability of a function, we define a canonisation operation on constraining functions, which preserves the set of im-

^{*} As the set of possible states of i is predefined by γ , it is not technically possible to add variables to it. The technical operation that we denote as “adding a variable” to the state of i in P_0 involves utilizing states of i that are not utilized by P_0 : Let $S_i^P \subseteq S_i$ be the set of states of i that are utilized in P (i.e. the union of its set of initial states \tilde{S}_i , with the image of the first coordinate of P_{0i}). We choose, as the set of states of i utilized by P , a subset of S_i that is in one-to-one correspondence with $\tilde{S}_i \times V$, where V is the set of possible values of the variable we wish to “add” to the state of i . As noted in Chapter 3, we assume that S_i is of large enough cardinality to allow for the existence of such a subset thereof.

[§] Ideality for solvability, in this sense, only reveals part of the whole picture, as it disregards the question of how fast can the responses be coordinated after an occurrence of \tilde{e} .

plementations. The canonical form of a constraining function will aid us in other aspects of the analysis of the timely-coordinated response problem as well, due to Remark 4.14. In order to define this canonical form, we consider δ as a weight function on the edges of a directed graph on I .

Definition 5.3. *Given an implementation-spec (I, δ) , we define the weighted directed graph of δ as $G_\delta \triangleq (I, E_\delta, \delta|_{E_\delta})$, where $E_\delta \triangleq \{(i, j) \in I^2 \mid \delta(i, j) \neq \infty\}$.*

Remark 5.4. *Let (I, δ) be an implementation-spec. By the above definition:*

1. *If $I = \{i, j\}$, then every $\bar{p} \in \mathcal{P}(G_\delta)$ is either of the form (i, j, i, j, \dots) or of the form (j, i, j, i, \dots) , for some $n \in \mathbb{N}$. (If $|I| > 2$, then $\mathcal{P}(G_\delta)$ is much richer.)*
2. $\forall \bar{p} \in \mathcal{P}(G_\delta) : L_{G_\delta}(\bar{p}) < \infty$.

Definition 5.5 (Canonical Form). *Let (I, δ) be an implementation-spec. We define the “canonical form” of δ as $\hat{\delta} \triangleq \delta_{G_\delta}$, the distance function of G_δ . By slight abuse of notation, we allow ourselves to write $\hat{\delta}$ instead of $\hat{\delta}|_{I^2}$ on some occasions below.*

Remark 5.6. *Let (I, δ) be an implementation-spec. By the above definition, $\hat{\delta}$ satisfies:*

- $\forall i \in I : \hat{\delta}(i, i) \in \{0, -\infty\}$. (Thus, by Remark 4.12, for implementable δ we obtain $\hat{\delta}|_{\{(i, i) \mid i \in I\}} = 0$.) Furthermore, $\hat{\delta}(i, i) = -\infty$ iff i is a vertex along a negative cycle in G_δ .
- Idempotence: $\hat{\delta} = \hat{\delta}$.
- Minimality: $\hat{\delta} \leq \delta$.
- Triangle inequality: $\forall i, j, k \in I : \hat{\delta}(i, k) \leq \hat{\delta}(i, j) + \hat{\delta}(j, k)$.
- Equivalence: $T(\delta) = T(\hat{\delta})$. (\subseteq : by the triangle inequality for path lengths. \supseteq : by minimality and by Remark 4.12 (monotonicity of T).)
- Monotonicity: Let $\delta' : I^2 \rightarrow \Delta$. If $\delta \leq \delta'$, then $\hat{\delta} \leq \hat{\delta}'$.

We are now ready to characterise the implementable functions from I^2 to Δ . The first part of the following lemma performs this task, while its second part shows that for every implementable δ , there exists an implementation that is minimal in every

coordinate — a result that gives us hope to find an optimal response logic for the timely-coordinated response problem for every δ .*

Lemma 5.7. *Let (I, δ) be an implementation-spec.*

1. δ is implementable iff $\hat{\delta}|_{\{i\} \times I}$ is bounded from below for every $i \in I$.
2. If δ is implementable, then $i \mapsto -\inf(\hat{\delta}|_{\{i\} \times I})$ is an implementation thereof, which is minimal in each coordinate.[§]

Proof. We first prove that if δ is implementable, then every $t \in T(\delta)$ satisfies $t(i) \geq -\inf(\hat{\delta}|_{\{i\} \times I})$. This implies the first direction (“ \Rightarrow ”) of the first part, and the inequality in the second part.

Assume that δ is implementable and let $t \in T(\delta)$ be an implementation thereof. By Remark 5.6 (equivalence), $t \in T(\hat{\delta})$ as well. Let $i \in I$. By definition of an implementation, we obtain

$$\forall j \in I \setminus \{i\} : \hat{\delta}(i, j) \geq t(j) - t(i) \geq 0 - t(i) = -t(i).$$

By Remark 5.6, $\hat{\delta}(i, i) = 0 \geq -t(i)$. Thus, we have $\hat{\delta}|_{\{i\} \times I} \geq -t(i)$. Taking the infimum over I of both sides of this inequality completes this part of the proof.

We now prove that if $\delta|_{\{i\} \times I}$ is bounded from below for every $i \in I$, then the function defined in the second part is indeed an implementation of δ . This completes the proof of both parts.

Define $t : I \rightarrow \mathbb{T}$ by $i \mapsto -\inf(\delta|_{\{i\} \times I}) < \infty$. By Remark 5.6, $\delta(i, i) \leq 0$, and therefore indeed $t \geq 0$. Let $(i, j) \in I^2$. Let $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$ s.t. $p_1 = j$. Define $p_0 \triangleq i$. Note that

$$\inf(\delta|_{\{i\} \times I}) \leq L_{G_\delta}((p_m)_{m=0}^n) = \delta(i, j) + L_{G_\delta}((p_m)_{m=1}^n).$$

* We do not wish, by any means, to imply that the existence of such a minimal implementation implies the existence of such an optimal response logic, nor even that in every run r of a protocol endowed with such a response logic, t_r is similar to this minimal implementation. We merely note that in the hypothetical absence of such a minimal implementation for some implementable δ , it would have been possible to show that no optimal response logic exists for the timely-coordinated response problem based on that δ . See also a discussion regarding Remark 6.36 below.

§ A quick glance at this formulation of the minimal implementation may raise a suspicion that perhaps it would have been more natural to define δ as the negation (in each coordinate) of the definition we have given. While it is indeed possible to define δ this way, and while doing so would have indeed given a more natural definition of the minimal implementation, it would have also required us to work with greatest path lengths instead of distances, with a reverse triangle inequality and with order-reversing monotonicity, which may somehow seem less natural.

By taking the infimum of both sides over all $\bar{p} \in \mathcal{P}(G_\delta)$ s.t. $p_1 = j$, we obtain $\inf(\delta|_{\{i\} \times I}) \leq \delta(i, j) + \inf(\delta|_{\{j\} \times I})$. Thus, we have $t(j) \leq t(i) + \delta(i, j)$, as required. \square

For the case in which I is finite, the first part of Lemma 5.7 yields the following, more tangible, implementation criterion.

Corollary 5.8. *Let (I, δ) be an implementation-spec s.t. $|I| < \infty$ and $\delta > -\infty$. δ is implementable iff G_δ contains no negative cycles.*

For completeness, we now prove a uniqueness property one may expect from the canonical form defined above, showing that the equivalence classes of implementable constraining functions, under the equivalence relation $\delta_1 \sim \delta_2 \Leftrightarrow T(\delta_1) = T(\delta_2)$, are in one-to-one, order-preserving, correspondence with canonical forms. At the heart of the proof of this property lies the following lemma.

Lemma 5.9. *Let (I, δ) be an implementation-spec s.t. δ is implementable, and let $\tilde{i}, \tilde{j} \in I$.*

1. *If $\hat{\delta}(\tilde{i}, \tilde{j}) < \infty$, then there exists an implementation $t \in T(\delta)$ satisfying $t(\tilde{j}) - t(\tilde{i}) = \hat{\delta}(\tilde{i}, \tilde{j})$.*
2. *If $\hat{\delta}(\tilde{i}, \tilde{j}) = \infty$, then for every $K \in \mathbb{T}$, there exists an implementation $t \in T(\delta)$ satisfying $t(\tilde{j}) - t(\tilde{i}) \geq K$.*

Proof. By Lemma 5.7, $\forall i \in I : \exists d_i \in \mathbb{T} : \hat{\delta}|_{\{i\} \times I} \geq -d_i$. (For the time being, we may choose $(-d_i)_{i \in I}$ to be the infima of the respective restrictions of $\hat{\delta}$.) We define $\delta' : I^2 \rightarrow \Delta \setminus \{\infty\}$ by

$$\forall i, j \in I^2 : \delta'(i, j) = \begin{cases} \delta(i, j) & \delta(i, j) < \infty \\ d_j & \delta(i, j) = \infty. \end{cases}$$

As $\delta' \leq \delta$, by monotonicity of T it is enough to find an implementation of δ' that satisfies the conditions of the lemma. By Remark 5.6 (minimality), it is enough find such an implementation for $\hat{\delta}'$.

We first show that $\forall i \in I : \hat{\delta}'|_{\{i\} \times I} \geq -d_i$. Let $i \in I$ and let $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_{\delta'})$ s.t. $p_1 = i$. Set $l = |\{k \in [n-1] \mid \delta(p_k, p_{k+1}) = \infty\}|$ — the number of “new” edges in \bar{p} , which do not exist in G_δ . We show, by induction on l , that $L_{G_{\delta'}}(\bar{p}) \geq -d_i$.

Base: If $l = 0$, then $L_{G_{\delta'}}(\bar{p}) = L_{G_{\delta}}(\bar{p}) \geq -d_i$.

Induction step: Assume $l \geq 1$. Let $k \in [n-1]$ be maximal such that (p_k, p_{k+1}) is a “new” edge (i.e. $\delta(p_k, p_{k+1}) = \infty$). By definition of δ' , we have $\delta'(p_k, p_{k+1}) = d_{p_{k+1}}$. Thus, by the induction hypothesis, we obtain

$$L_{G_{\delta'}}(\bar{p}) = L_{G_{\delta'}}((p_m)_{m=1}^k) + \delta'(p_k, p_{k+1}) + L_{G_{\delta}}((p_m)_{m=k+1}^n) \geq -d_i + d_{p_{k+1}} - d_{p_{k+1}} = -d_i.$$

and the proof by induction is complete. In particular, we conclude that $\widehat{\delta'} > -\infty$, and by definition, also $\widehat{\delta'} \leq \delta' < \infty$.

We claim that $t \triangleq d_i + \widehat{\delta'}(\tilde{i}, \cdot) \geq 0$ is an implementation of $\widehat{\delta'}$. Indeed, for every $(j, k) \in I^2$, by Remark 5.6 (triangle inequality), we have

$$t(k) = d_i + \widehat{\delta'}(\tilde{i}, k) \leq d_i + \widehat{\delta'}(\tilde{i}, j) + \widehat{\delta'}(j, k) = t(j) + \widehat{\delta'}(j, k).$$

If $\widehat{\delta'}(\tilde{i}, \tilde{j}) < \infty$, we define $K \triangleq \widehat{\delta'}(\tilde{i}, \tilde{j})$; otherwise, let $K \in \mathbb{T}$ be arbitrarily large as in the conditions of the lemma. As δ' is implementable, by Remark 5.6 we obtain $\widehat{\delta'}(\tilde{i}, \tilde{i}) = 0$. Therefore,

$$t(\tilde{j}) - t(\tilde{i}) = (d_i + \widehat{\delta'}(\tilde{i}, \tilde{j})) - (d_i + \widehat{\delta'}(\tilde{i}, \tilde{i})) = \widehat{\delta'}(\tilde{i}, \tilde{j}).$$

Thus, if $\widehat{\delta'}(\tilde{i}, \tilde{j}) \geq K$, then the proof is complete. (For the case in which $\widehat{\delta'}(\tilde{i}, \tilde{j}) < \infty$, we obtain $\widehat{\delta'}(\tilde{i}, \tilde{j}) \leq K$ by Remark 5.6 (monotonicity), since $\delta' \leq \delta$.)

Otherwise, set $d \triangleq K - \widehat{\delta'}(\tilde{i}, \tilde{j}) > 0$, and define $d'_i \triangleq d_i + d > d_i$, for every $i \in I$. Therefore, $-d'_i < -d_i \leq \widehat{\delta}|_{\{i\} \times I}$ for every $i \in I$. Denote by δ'' the function constructed from δ in the same way in which δ' was constructed from it, but using the lower bounds $(-d'_i)_{i \in I}$ rather than $(-d_i)_{i \in I}$. As explained above, in order to complete the proof it is enough to show that $\widehat{\delta''}(\tilde{i}, \tilde{j}) \geq K$. Let $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_{\delta''})$ s.t. $p_1 = \tilde{i}$ and $p_n = \tilde{j}$. If $\forall k \in [n-1] : \delta(p_k, p_{k+1}) \neq \infty$, then $L_{G_{\delta''}}(\bar{p}) = L_{G_{\delta}}(\bar{p}) \geq \widehat{\delta}(\tilde{i}, \tilde{j}) \geq K$. Otherwise,

$$\begin{aligned} L_{G_{\delta''}}(\bar{p}) &= && \text{by definitions of } \delta' \text{ and } \delta'' \\ &= L_{G_{\delta'}}(\bar{p}) + d \cdot |\{k \in [n-1] \mid \delta(p_k, p_{k+1}) = \infty\}| \geq && \text{as this set is non-empty} \\ &\geq L_{G_{\delta'}}(\bar{p}) + d \geq && \text{by definition of } \widehat{\delta'} \\ &\geq \widehat{\delta'}(\tilde{i}, \tilde{j}) + d = && \text{by definition of } K \\ &= K. \end{aligned}$$

Either way, the proof is complete. \square

While unimplementable functions whose canonical forms differ may exist (due to I not necessarily being finite, and due to the fact that G_δ needs not necessarily be strongly connected), we now conclude, using Lemma 5.9, that for implementable functions, the map $\hat{\delta} \mapsto T(\delta)$ from the canonical form of an implementable function δ to the set of implementations of δ is a well-defined, order-preserving, monomorphism.

Corollary 5.10. *Let I be a set and let $\delta_1, \delta_2 : I^2 \rightarrow \Delta$ s.t. δ_1 is implementable. $\hat{\delta}_1 \leq \hat{\delta}_2$ iff $T(\delta_1) \subseteq T(\delta_2)$.*

Proof. \Rightarrow : Assume that $\hat{\delta}_1 \leq \hat{\delta}_2$. By monotonicity of T and by Remark 5.6 (equivalence), we have $T(\delta_1) = T(\hat{\delta}_1) \subseteq T(\hat{\delta}_2) = T(\delta_2)$.

\Leftarrow : Assume that $\hat{\delta}_1 \not\leq \hat{\delta}_2$. Thus, there exist $\tilde{i}, \tilde{j} \in I$ s.t. $\hat{\delta}_1(\tilde{i}, \tilde{j}) > \hat{\delta}_2(\tilde{i}, \tilde{j})$. If $\hat{\delta}_1(\tilde{i}, \tilde{j}) < \infty$, then by Lemma 5.9 there exists $t \in T(\delta_1)$ s.t. $t(\tilde{j}) - t(\tilde{i}) = \hat{\delta}_1(\tilde{i}, \tilde{j}) > \hat{\delta}_2(\tilde{i}, \tilde{j})$, and thus $t \in T(\delta_1) \setminus T(\delta_2)$, and the proof is complete.

If $\hat{\delta}_1(\tilde{i}, \tilde{j}) = \infty$, then $\hat{\delta}_2(\tilde{i}, \tilde{j}) < \infty$ and thus there exists $K \in \mathbb{T}$ s.t. $K > \hat{\delta}_2(\tilde{i}, \tilde{j})$. Similarly to the proof of the previous case, by Lemma 5.9 there exists $t \in T(\delta_1)$ s.t. $t(\tilde{j}) - t(\tilde{i}) \geq K > \hat{\delta}_2(\tilde{i}, \tilde{j})$. Once again, we obtain that $t \in T(\delta_1) \setminus T(\delta_2)$, and the proof is complete. \square

Corollary 5.11. *Let I be a set and let $\delta_1, \delta_2 : I^2 \rightarrow \Delta$ s.t. at least one of them is implementable. $\hat{\delta}_1 = \hat{\delta}_2$ iff $T(\delta_1) = T(\delta_2)$.*

Proof. \Rightarrow : Assume that $\hat{\delta}_1 = \hat{\delta}_2$. By applying Remark 5.6 (equivalence) twice, we obtain $T(\delta_1) = T(\hat{\delta}_1) = T(\hat{\delta}_2) = T(\delta_2)$.

\Leftarrow : Assume that $T(\delta_1) = T(\delta_2)$. Thus, since at least one of δ_1, δ_2 is implementable, they both are. To complete the proof, we apply Corollary 5.10 to $T(\delta_1) \subseteq T(\delta_2)$ and to $T(\delta_2) \subseteq T(\delta_1)$. \square

The above discussion gives rise to two alternative definitions (or rather, characterisations) of the canonical form of implementable functions: The first one, non-constructive in nature, justifies the name of the minimality property stated in Remark 5.6 and stems from this property when combined with Corollary 5.11. The second one, which constructively defines the inverse of the order-preserving monomorphism $\hat{\delta} \mapsto T(\delta)$, stems directly from Lemma 5.9.

Corollary 5.12. *Let (I, δ) be an implementation-spec s.t. δ is implementable.*

1. $\hat{\delta} = \min\{\delta' \in \Delta^{(I^2)} \mid T(\delta') = T(\delta)\}$. (In particular, there exists a function in this set, which is minimal in each coordinate, although this may be proven directly by means of a simpler argument.)
2. $\forall i, j \in I : \hat{\delta}(i, j) = \max\{t(j) - t(i) \mid t \in T(\delta)\}$.

Remark 5.13. Implementability of δ is not required in Corollary 5.12 if $|I| < \infty$ and if G_δ is strongly connected. Indeed, under such conditions, if δ is unimplementable, then $\hat{\delta} \equiv -\infty$, which coincides with the function obtained in both parts of this corollary, when they are applied to any unimplementable δ .

By applying Claim 5.1 to the previous three corollaries, we obtain similar results regarding the map $\hat{\delta} \mapsto TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ from the canonical form of an implementable function δ to the timely-coordinated response problem that δ defines with respect to a fixed external input. We conclude this chapter by formulating these results.

Corollary 5.14. Let γ be a context, let $I \subseteq \mathbb{I}_\gamma$ and let $\tilde{e} \in \tilde{E}_\gamma$ s.t. $TCR_\gamma\langle\tilde{e}, I, 0\rangle$ is solvable.*

- Let $\delta_1, \delta_2 : I^2 \rightarrow \Delta$.
 1. If δ_1 is implementable, then: $\hat{\delta}_1 \leq \hat{\delta}_2$ iff $TCR_\gamma\langle\tilde{e}, I, \delta_1\rangle \subseteq TCR_\gamma\langle\tilde{e}, I, \delta_2\rangle$.
 2. If either δ_1 or δ_2 are implementable, then: $\hat{\delta}_1 = \hat{\delta}_2$ iff $TCR_\gamma\langle\tilde{e}, I, \delta_1\rangle = TCR_\gamma\langle\tilde{e}, I, \delta_2\rangle$.
- Let $\delta : I^2 \rightarrow \Delta$ be implementable. (Once again, implementability of δ is not required for this part if $|I| < \infty$ and if G_δ is strongly connected.)
 1. $\hat{\delta} = \min\{\delta' \in \Delta^{(I^2)} \mid TCR_\gamma\langle\tilde{e}, I, \delta'\rangle = TCR_\gamma\langle\tilde{e}, I, \delta\rangle\}$.
 2. $\forall i, j \in I : \hat{\delta}(i, j) = \max\{t_r(j) - t_r(i) \mid r \in \cup_{P \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle} R_\gamma(P)\}$.

Readers who found our choice from the previous chapter, to formally associate a coordinated response problem with the set of solutions thereof philosophically troubling, may find some justification for this choice in the first part of Corollary 5.14. This part essentially shows that in order to accept our choice, at least when $TCR_\gamma\langle\tilde{e}, I, 0\rangle$ is solvable and for implementable δ , it is enough to accept that $TCR_\gamma\langle\tilde{e}, I, \hat{\delta}\rangle$ is the same problem as $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$.

* As noted above regarding Remark 5.2, Corollary 9.7 shows that under certain conditions, the 0 function in Corollary 5.14 may be replaced with a variety of other functions.

Chapter 6

The Syncausality Approach

In this chapter, we analyze the timely-coordinated response problem using tools developed by Ben-Zvi and Moses[6, 5], and generalize some previous results obtained by them[6, 5, 7, 8] using these tools. The proofs that we give in this chapter, unlike the proofs in [6, 5, 7, 8], do not explicitly use the concept of knowledge. We choose to phrase our proofs in this way in order to emphasize the difference between the approach taken in this chapter and that of the next one.

6.1 Background

This section surveys previous definitions and results from [6, 5, 7, 8].

6.1.1 Partial Orders on the Set of Agent-Time Pairs

Recall that we work in a context consisting of a set of agents \mathbb{I}_γ that communicate with each other solely via message passing, and that their communication channels are modeled by the edges of the directed graph $\mathcal{G}_\gamma = (\mathbb{I}_\gamma, N_\gamma, b_\gamma)$, each of which being weighted according to the maximum delivery time of a message along it.

Ben-Zvi and Moses[6, 5] define two partial order relations on the set of agent-time pairs $\mathbb{I}_\gamma \times \mathbb{T}$. The first, called “syncausality” (short for synchronous causality), is a synchronous counterpart to Lamport’s “happened-before” causality relation[17], and similarly aims to capture information flow. Intuitively, if $(i, t) \xrightarrow[r]{\gamma} (j, t')$, then in some sense, j potentially has, at t' in r , information regarding the state of i at t in r . (In a full-information protocol, this intuition can be made more concrete: the state of i at t in r can be deduced with absolute certainty from the state of j at t'

in r .)

Definition 6.1 (Syncausality). *Let γ be a context and let $r \in \mathcal{R}_\gamma$. The “Syncausality” relation $\overset{\gamma}{\rightsquigarrow}_r$ is the minimal partial order relation on $\mathbb{I}_\gamma \times \mathbb{T}$ satisfying (i.e. the transitive closure of)*

- *Locality: $\forall i \in \mathbb{I}_\gamma, t, t' \in \mathbb{T} : t' \geq t \Rightarrow (i, t) \overset{\gamma}{\rightsquigarrow}_r (i, t')$.*
- *Message delivery: If, in r , a message is sent from $i \in \mathbb{I}_\gamma$ at $t \in \mathbb{T}$ and delivered to $j \in \mathbb{I}_\gamma$ at $t' > t$, then $(i, t) \overset{\gamma}{\rightsquigarrow}_r (j, t')$.*
- *Delivery guarantee: $(i, t) \overset{\gamma}{\rightsquigarrow}_r (j, t + b_\gamma(i, j))$, for every $(i, j) \in N_\gamma$ s.t. $b_\gamma(i, j) < \infty$, and for every $t \in \mathbb{T}$.*

The syncausality relation is a refinement of Lamport’s “happened-before” causality relation[17], which is defined similarly, with the only difference being the absence of the delivery guarantee property. At first sight, this property may seem redundant due to the message delivery property. Indeed, the bound guarantee property is of importance only if a message is *not* sent from i to j at t . Intuitively, j has a guarantee that i did not send it a message at t , only when the worst-case delivery time for such a message has elapsed, i.e. at $t + b_\gamma(i, j)$. Passing information by *not* sending a message was first studied by Lamport[18], who called such unsent messages “null messages”.

The second partial order relation on $\mathbb{I}_\gamma \times \mathbb{T}$, called “bound guarantee”, aims to capture guaranteed information flow, and thus has no asynchronous counterpart. Intuitively, if $(i, t) \overset{\gamma}{\dashrightarrow} (j, t')$, then in some sense, it is not only that j potentially has, at t' in any run r , information regarding i at t in r , but also that i has some guarantee at t in r that such information has the potential to reach j by t' in r at the latest. (In a full-information protocol, this means that from the state of i at t in r , it can be deduced (with absolute certainty) that this state may be deduced from the state of j at t' in r .)

Definition 6.2 (Bound Guarantee). *Let γ be a context. The “Bound Guarantee” relation $\overset{\gamma}{\dashrightarrow}$ is the minimal partial order relation on $\mathbb{I}_\gamma \times \mathbb{T}$ satisfying:*

- *Locality: $\forall i \in \mathbb{I}_\gamma, t, t' \in \mathbb{T} : t' \geq t \Rightarrow (i, t) \overset{\gamma}{\dashrightarrow} (i, t')$.*
- *Delivery guarantee: $(i, t) \overset{\gamma}{\dashrightarrow} (j, t + b_\gamma(i, j))$, for every $(i, j) \in N_\gamma$ s.t. $b_\gamma(i, j) < \infty$, and for every $t \in \mathbb{T}$.*

Remark 6.3. *As would be expected by the intuitive descriptions of both relations above, the bound guarantee relation is a subrelation of any syncausality relation: If $(i, t) \dashrightarrow^\gamma (j, t')$, then $\forall r \in \mathcal{R}_\gamma : (i, t) \xrightarrow[r]{\gamma} (j, t')$.*

6.1.2 Additional Notation

We now introduce some novel notation, which aims to capture the flow of information regarding the occurrence of events, and the flow of information which may affect the occurrence of an event. This notation will both aid us in more succinctly presenting some previous results of Ben-Zvi and Moses in the next subsection, and in presenting our results thereafter.

Definition 6.4. *Let γ be a context and let $r \in \mathcal{R}_\gamma$.*

1. *Given an event $e \in E(r)$ and an agent-time pair $(i, t) \in \mathbb{I}_\gamma \times \mathbb{T}$, we write $e \xrightarrow[r]{\gamma} (i, t)$ (resp. $e \dashrightarrow^\gamma (i, t)$) if $(i_e, t_e) \xrightarrow[r]{\gamma} (i, t)$ (resp. $(i_e, t_e) \dashrightarrow^\gamma (i, t)$), where by i_e we denote the immediate observer of e . (Recall from Chapter 3, that t_e is the occurrence time of e .)*
2. *Given two events $e, e' \in E(r)$, we write $e \xrightarrow[r]{\gamma} e'$ (resp. $e \dashrightarrow^\gamma e'$) if either $e = e'$ or e' is a delivery of a message sent by an agent $i \in \mathbb{I}_\gamma$ at time $t \in \mathbb{T}$ s.t. $e \xrightarrow[r]{\gamma} (i, t)$ (resp. $e \dashrightarrow^\gamma (i, t)$).*

Once again, in a full-information protocol, some of the implications of these definitions become very concrete, e.g. $e \xrightarrow[r]{\gamma} (i, t)$ guarantees that the occurrence of e may be deduced from the state of i at t . Similarly, $e \xrightarrow[r]{\gamma} e'$ guarantees that if e' is a message event, then the occurrence of e may be deduced from the contents of the message associated with e' . Moreover, if $e \xrightarrow[r]{\gamma} e'$ does *not* hold, then the occurrence of e' does not depend, in a sense, on the occurrence of e . We make this last observation precise in Corollary 6.30 below.

6.1.3 Previous Results

In this subsection, we survey the coordinated response problems defined and studied by Ben-Zvi and Moses in [6, 5, 7, 8], and their results for these problems in discrete-time models. (The only coordinated response problem from [6, 5, 7, 8] that we do not survey in this subsection, namely “general ordered response”, is discussed

in Chapter 10.) We reformulate these problems, results, and the associated definitions to match our notation, and to make use of our coordinated-response-theoretic definitions.

While surveying all these coordinated response problems, and while remarking, by defining an appropriate δ function, that the timely-coordinated response problem extends each and every one of them (and also extends general ordered response), one property, which is common to all these δ functions, should be spelled out explicitly: $\hat{\delta}$ is antisymmetric on each strongly-connected component of G_δ .^{*} As we will see during this work, the absence of this property in the timely-coordinated response problem introduces a significant amount of complexity, both technically, and conceptually.

The first, most-basic coordinated response problem defined in [6, 5] is that of ordered response.

Definition 6.5 (Ordered Response). *Given a context γ , an external input $\tilde{e} \in \tilde{E}_\gamma$, $n \in \mathbb{N}$ and agents $\bar{i} = (i_m)_{m=1}^n \in \mathbb{I}_\gamma^n$, define the “ordered response” problem $OR_\gamma\langle\tilde{e}, \bar{i}\rangle \subseteq ER_\gamma\langle\tilde{e}, \{i_m\}_{m=1}^n\rangle$ as the set of all eventual-response protocols P satisfying $t_r(i_{m+1}) \geq t_r(i_m)$ for every $m \in [n-1]$ and for every triggered run $r \in R_\gamma^{\tilde{e}}(P)$.*

Remark 6.6. $OR_\gamma\langle\tilde{e}, (i_m)_{m=1}^n\rangle = TCR_\gamma\langle\tilde{e}, \{i_m\}_{m=1}^n, \delta\rangle$, for

$$\delta(i_k, i_l) \triangleq \begin{cases} 0 & k = l + 1 \\ \infty & \text{otherwise.} \end{cases}$$

Ben-Zvi and Moses analyze this problem using a structure they call “centipede”.

Definition 6.7 (Centipede — see Figure 1). *Given a context γ , a run $r \in \mathcal{R}_\gamma$, an external input $\tilde{e} \in \tilde{E}_\gamma$, $n \in \mathbb{N}$, and agents $\bar{i} \in \mathbb{I}_\gamma^n$, with matching times $\bar{t} = (t_m)_{m=1}^n \in \mathbb{T}^n$, we call an n -tuple of ND events $\bar{e} = (e_m)_{m=1}^n \in ND_\gamma(r)^n$ an “ \tilde{e} -centipede” for \bar{i} by \bar{t} if the following hold:*

- $\tilde{e} \xrightarrow{\gamma}_r e_1$ and $\forall m \in [n-1] : e_m \xrightarrow{\gamma}_r e_{m+1}$.
- $\forall m \in [n] : e_m \xrightarrow{\gamma} (i_m, t_m)$.

^{*} It is interesting to note, though, that some instances of the timely-coordinated response problem, while having this property, are not instances of any of the problems defined and studied by Ben-Zvi and Moses. We analyze such instances in the second part of Corollary 9.5.

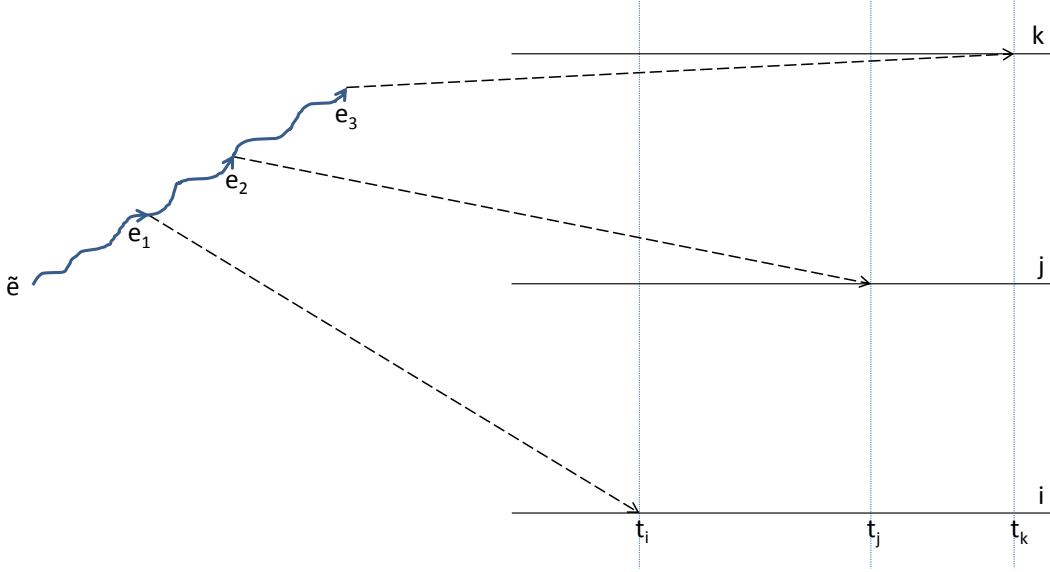


Figure 1: (e_1, e_2, e_3) is an \tilde{e} -centipede for $\{i, j, k\}$ by (t_i, t_j, t_k) .

Given a time $t \in \mathbb{T}$, we call an n -tuple of ND events $\bar{e} \in ND_\gamma(r)^n$ an “ \tilde{e} -centipede” for \bar{i} by t , if it is an \tilde{e} -centipede for \bar{i} by $(t)^n$.

Theorem 6.8 (Centipede). *In a discrete-time model, let γ be a context, let $n \in \mathbb{N}$, let $\bar{i} \in \mathbb{I}_\gamma^n$ and let $\tilde{e} \in \tilde{E}_\gamma$.*

1. *Let $P \in OR_\gamma\langle\tilde{e}, \bar{i}\rangle$. Each $r \in R_\gamma^{\tilde{e}}(P)$ contains an \tilde{e} -centipede for \bar{i} by $t_r(i_n)$.*
2. *In a shared-clock model, an optimal response logic for solving $OR_\gamma\langle\tilde{e}, \bar{i}\rangle$ is, for every i_m : “respond at the earliest time by which an \tilde{e} -centipede for $(i_k)_{k=1}^m$ exists.”*

Ben-Zvi and Moses prove the first part of Theorem 6.8 in two stages: First, reducing to a response-recalling protocol (a protocol in which the set of responses of an agent up until time t may be deduced from its state at t), they show that under the conditions of that part of the theorem, at $t_r(i_n)$ it holds that $K_{i_n} \cdots K_{i_1} \tilde{e}$ (where K_i means “ i knows that...” — this will be formally defined in Chapter 7). Second, conceptually following the path of Chandy and Misra[9], they deduce the existence of the required centipede from this nested-knowledge formula.

Regarding the second part of Theorem 6.8, it should be noted that Ben-Zvi and Moses do not define a notion of optimality, but rather show that a full-information protocol with the given response logic solves $OR_\gamma\langle\tilde{e}, \bar{i}\rangle$. (Optimality, under our definition, may be derived from the combination of the two parts of Theorem 6.8,

which Ben-Zvi and Moses prove as separate theorems.) A similar note holds for the second part of Theorem 6.12 below.

The second problem presented in [6, 5] is the following variant of the firing squad problem [21, 10].

Definition 6.9 (Simultaneous Response). *Given a context γ , an external input $\tilde{e} \in \tilde{E}_\gamma$ and set of agents $I \subseteq \mathbb{I}_\gamma$, define the “simultaneous response” problem $SR_\gamma\langle\tilde{e}, I\rangle \subseteq ER_\gamma\langle\tilde{e}, I\rangle$ as the set of all eventual-response protocols P for which t_r is a constant function for each run $r \in R_\gamma(P)$. We denote, in this case, the constant value of t_r by t_r .*

Remark 6.10. $SR_\gamma\langle\tilde{e}, I\rangle = TCR_\gamma\langle\tilde{e}, I, 0\rangle$.

Ben-Zvi and Moses analyze this problem using a structure they call “broom”.

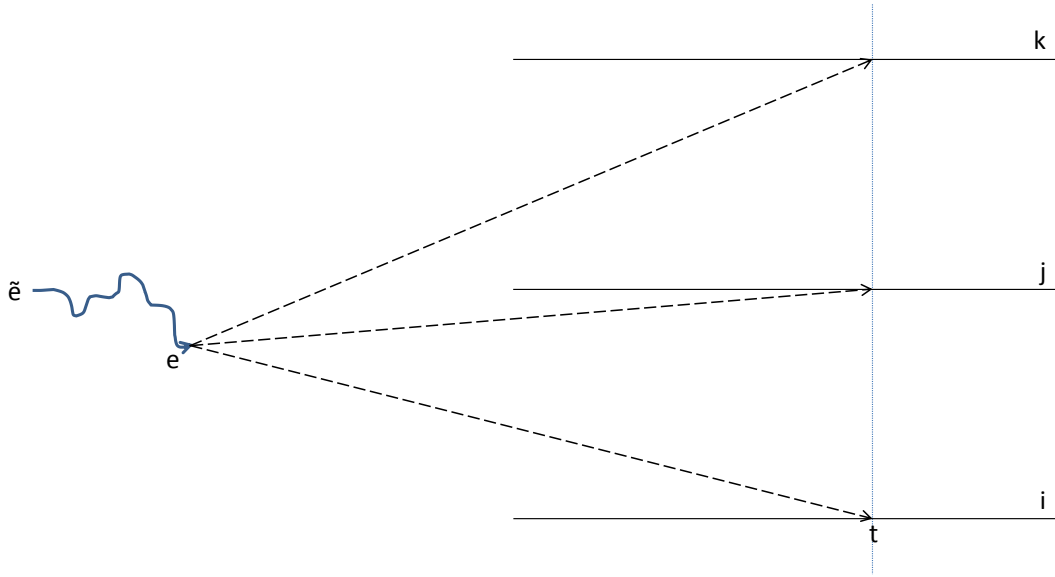


Figure 2: (e_1, e_2, e_3) is an \tilde{e} -broom for $\{i, j, k\}$ by t .

Definition 6.11 (Broom — see Figure 2). *Given a context γ , a run $r \in \mathcal{R}_\gamma$, an external input $\tilde{e} \in \tilde{E}_\gamma$ and a set of agents $I \subseteq \mathbb{I}_\gamma$ with matching times $\bar{t} = (t_i)_{i \in I} \in \mathbb{T}^I$, we call an ND event $e \in ND_\gamma(r)$ an “ \tilde{e} -broom” for I by \bar{t} if the following hold:*

- $\tilde{e} \xrightarrow[\gamma]{r} e$
- $\forall i \in I : e \xrightarrow[\gamma]{} (i, t_i)$. We call $\{(i, t_i)\}_{i \in I}$ the set of “end nodes” of this broom, and call $\max\{\bar{t}\}$ the “horizon” of this broom.

Given a time $t \in \mathbb{T}$, we call an ND event $e \in ND_\gamma(r)$ an “ \tilde{e} -broom” for I by t , if it is an \tilde{e} -broom for I by $(t)^I$.

Theorem 6.12 (Broom). *In a discrete-time model,* let γ be a context, let $I \subseteq \mathbb{I}_\gamma$ be finite, and let $\tilde{e} \in \tilde{E}_\gamma$.*

- *Let $P \in SR_\gamma\langle\tilde{e}, I\rangle$. Each $r \in R_\gamma^\tilde{e}(P)$ contains an \tilde{e} -broom for I by t_r .*
- *In a shared-clock model, an optimal response logic for solving $SR_\gamma\langle\tilde{e}, I\rangle$ is, for every $i \in I$: “respond at the earliest time by which an \tilde{e} -broom for I exists.”*

Ben-Zvi and Moses prove the first part of Theorem 6.12 by reducing to a response-recalling protocol, showing that under the conditions of that part of the theorem, \tilde{e} is common knowledge among all agents in I at t_r , and then using a reduction to the first part of Theorem 6.8. We give a direct proof of a slight generalization of Theorem 6.12 later in this work (see Theorem 9.12).

The third and last problem presented in [5], is the following generalization of both ordered response and simultaneous response.

Definition 6.13 (Ordered Joint Response). *Given a context γ , an external input $\tilde{e} \in \tilde{E}_\gamma$, $n \in \mathbb{N}$ and pairwise-disjoint non-empty sets of agents $\bar{I} = (I_m)_{m=1}^n \in (2^{\mathbb{I}_\gamma})^n$, define the “ordered joint response” problem $OJR_\gamma\langle\tilde{e}, \bar{I}\rangle \subseteq ER_\gamma\langle\tilde{e}, \cup_{m=1}^n I_m\rangle$ as the set of all eventual-response protocols P satisfying, for every run $r \in R_\gamma^\tilde{e}(P)$:*

1. $t_r|_{I_m}$ is a constant function, for each $m \in [n]$. We denote its value by $t_{r,m}$.
2. $\forall m \in [n-1] : t_{r,m+1} \geq t_{r,m}$.

Remark 6.14. $OJR_\gamma\langle\tilde{e}, \bar{I}\rangle = TCR_\gamma\langle\tilde{e}, \cup_{m=1}^n I_m, \delta\rangle$, for

$$\delta(i, j) \triangleq \begin{cases} 0 & \exists k, l \in [n] : i \in I_k \ \& \ j \in I_l \ \& \ k \geq l \\ \infty & \text{otherwise.} \end{cases}$$

Ben-Zvi and Moses analyze this problem using a structure they call “centibroom”, which generalizes both a centipede and a broom.

Definition 6.15 (Centibroom — see Figure 3). *Given a context γ , a run $r \in \mathcal{R}_\gamma$, an external input $\tilde{e} \in \tilde{E}_\gamma$, $n \in \mathbb{N}$, and pairwise-disjoint non-empty sets of agents $\bar{I} \in (2^{\mathbb{I}_\gamma})^n$ with matching times $\bar{t} \in \mathbb{T}^I$, where $I \triangleq \cup_{m=1}^n I_m$, we call an n -tuple of ND events $\bar{e} \in ND_\gamma(r)^n$ an “ \tilde{e} -centibroom” for \bar{I} by \bar{t} if the following hold:*

* This is a key requirement here.

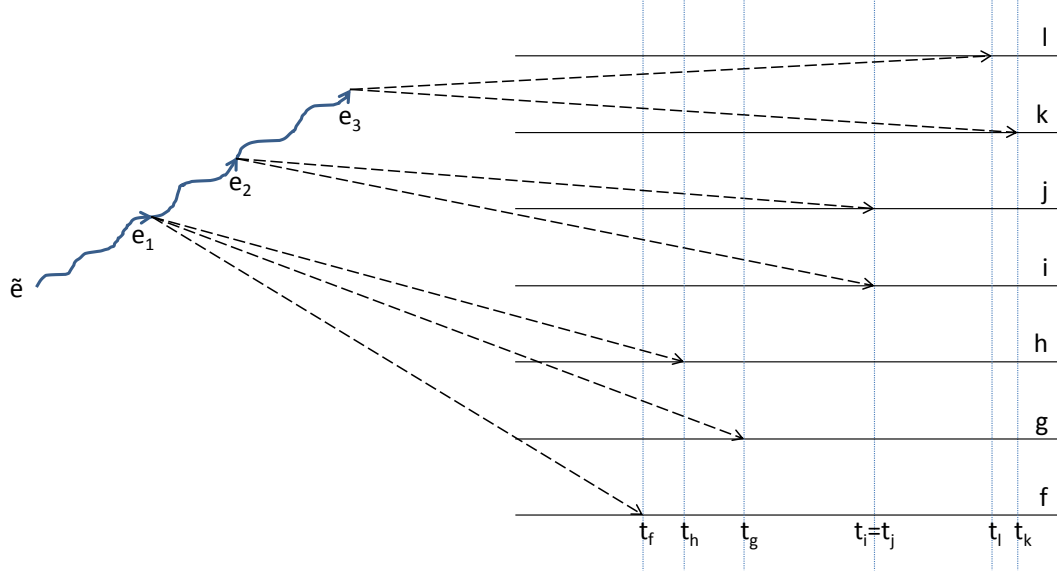


Figure 3: (e_1, e_2, e_3) is an \tilde{e} -centibroom for $(\{f, g, h\}, \{i, j\}, \{k, l\})$ by $(t_f, t_g, t_h, t_i, t_j, t_k, t_l)$

- $\tilde{e} \xrightarrow[r]{\gamma} e_1$ and $\forall m \in [n-1] : e_m \xrightarrow[r]{\gamma} e_{m+1}$.
- $\forall m \in [n], i \in I_m : e_m \xrightarrow{\gamma} (i, t_i)$. We call $\{(i, t_i)\}_{i \in I}$ the set of “end nodes” of this centibroom, and call $\max\{t\}$ the “horizon” of this centibroom.

Given a time $t \in \mathbb{T}$, we call an n -tuple of ND events $\bar{e} \in ND_\gamma(r)^n$ an “ \tilde{e} -centibroom” for \bar{I} by t , if it is an \tilde{e} -centibroom for \bar{I} by $(t)^n$.

Theorem 6.16 (Centibroom). *In a discrete-time model,* let γ be a context, let $n \in \mathbb{N}$, let $\bar{I} \in (2^{\mathbb{I}_\gamma})^n$ be pairwise-disjoint non-empty finite sets of agents, let $\tilde{e} \in \tilde{E}_\gamma$ and let $P \in OJR_\gamma\langle \tilde{e}, \bar{I} \rangle$. Each $r \in R_\gamma^{\tilde{e}}(P)$ contains an \tilde{e} -centibroom for \bar{I} by $t_{r,n}$.*

Ben-Zvi and Moses prove Theorem 6.16 by reducing to a response-recalling protocol, showing that under the conditions of that part of the theorem, at $t_{r,n}$ it holds that $C_{I_n} \cdots C_{I_1} \tilde{e}$ (where C_J means “it is common knowledge among J that...” — this will be formally defined in Chapter 7), and then using a reduction to the first part of Theorem 6.8.

In [8], the following respective generalizations of ordered response and simultaneous response were introduced: (once again, we rephrase them to match the definitions and notation introduced in this work.)

* Once again, this is a key requirement here.

Definition 6.17 (Weakly-Timed Response). *Given a context γ , an external input $\tilde{e} \in \tilde{E}_\gamma$, $n \in \mathbb{N}$, agents $\bar{i} \in \mathbb{I}_\gamma^n$ and finite time-differences $\bar{\varepsilon} = (\varepsilon_m)_{m=1}^{n-1} \in (\Delta \setminus \{-\infty, \infty\})^{n-1}$, define the “weakly-timed response” problem $WTR_\gamma\langle\tilde{e}, \bar{i}, \bar{\varepsilon}\rangle \subseteq ER_\gamma\langle\tilde{e}, \{i_m\}_{m=1}^n\rangle$ as the set of all eventual-response protocols P satisfying $t_r(i_{m+1}) \geq t_r(i_m) + \varepsilon_m$ for every $m \in [n-1]$ and for every triggered run $r \in R_\gamma^\tilde{e}(P)$.*

Remark 6.18. $WTR_\gamma\langle\tilde{e}, (i_m)_{m=1}^n, \bar{\varepsilon}\rangle = TCR_\gamma\langle\tilde{e}, \{i_m\}_{m=1}^n, \delta\rangle$, for

$$\delta(i_k, i_l) \triangleq \begin{cases} -\varepsilon_l & k = l + 1 \\ \infty & \text{otherwise.} \end{cases}$$

Definition 6.19 (Tightly-Timed Response). *Given a context γ , an external input $\tilde{e} \in \tilde{E}_\gamma$ and a set of agents $I \subseteq \mathbb{I}_\gamma$ with matching times $\bar{t} \in \mathbb{T}^I$, define the “simultaneous response” problem $TTR_\gamma\langle\tilde{e}, I, \bar{t}\rangle \subseteq ER_\gamma\langle\tilde{e}, I\rangle$ as the set of all eventual-response protocols P satisfying $t_r(i) - t_r(j) = t_i - t_j$ for every $i, j \in I$ and every run $r \in R_\gamma(P)$.*

Remark 6.20. $TTR_\gamma\langle\tilde{e}, I, \bar{t}\rangle = TCR_\gamma\langle\tilde{e}, I, \delta\rangle$, for $\delta(i, j) \triangleq t_j - t_i$.

Ben-Zvi and Moses present the following theorems in [8], and prove them along the lines of their proofs of the first part of Theorem 6.8 and the first part of Theorem 6.12, respectively.

Theorem 6.21 (Uneven Centipede). *In a discrete-time model, let γ be a context, let $n \in \mathbb{N}$, let $\bar{i} \in \mathbb{I}_\gamma^n$, let $\bar{\varepsilon} \in (\Delta \setminus \{-\infty, \infty\})^{n-1}$, let $\tilde{e} \in \tilde{E}_\gamma$ and let $P \in OR_\gamma\langle\tilde{e}, \bar{i}\rangle$. Each $r \in R_\gamma^\tilde{e}(P)$ contains an \tilde{e} -centipede for \bar{i} by $(t_r(i_n) - \sum_{k=m}^{n-1} \varepsilon_k)_{m=1}^n$.*

Theorem 6.22 (Uneven Broom). *In a discrete-time model,* let γ be a context, let $I \subseteq \mathbb{I}_\gamma$ be finite, let $\bar{t} \in \mathbb{T}^I$, let $\tilde{e} \in \tilde{E}_\gamma$ and let $P \in SR_\gamma\langle\tilde{e}, I\rangle$. Each $r \in R_\gamma^\tilde{e}(P)$ contains an \tilde{e} -broom for I by $(t_r(i))_{i \in I}$.*

6.2 Adapting Some Machinery

Before approaching the timely-coordinated response problem using the definitions surveyed in the previous section, we adapt some of the machinery used by Ben-Zvi and Moses to obtain the results surveyed therein. In order to do so, we introduce, yet again, some additional novel notation and definitions.

* Yet again, this is a key requirement here.

Definition 6.23. Let γ be a context and let $r \in \mathcal{R}_\gamma$.

- Given a time $t \in \mathbb{T}$, we denote the set of all ND events occurring in r no later than t by $ND_\gamma(r, t) \triangleq \{e \in ND_\gamma(r) \mid t_e \leq t\}$.
- Given an agent-time pair $\theta \in \mathbb{I}_\gamma \times \mathbb{T}$ (resp. an event $\theta \in E(r)$), we define the “ND past” of θ in r as $PND_\gamma^r(\theta) \triangleq \{e \in ND_\gamma(r) \mid e \xrightarrow[r]{\gamma} \theta\}$. Note that $PND_\gamma^r(\theta) \subseteq ND_\gamma(r, t)$, where $\theta = (i, t)$ (resp. $t = t_\theta$). For an agent-time pair (i, t) , we sometimes write $PND_\gamma^r(i, t)$ instead of $PND_\gamma^r((i, t))$, for readability.

Definition 6.24 (Subruns). Given a context γ , a protocol $P \in \mathbb{P}_\gamma$, a time $t \in \mathbb{T}$ and runs $r, r' \in R_\gamma(P)$, we call r' a “ t -subrun” of r , and write $r' \stackrel{t}{\subseteq} r$, if the initial states used for all agents in r' and in r are the same, and if $ND_\gamma(r', t) \subseteq ND_\gamma(r, t)$. (We omit P and γ from this notation for readability, as they will be clear from the discussion.) For fixed P and t , we note that $\stackrel{t}{\subseteq}$ is a quasi-order relation on $R_\gamma(P)$, in which two runs are in the same equivalence class iff they are indistinguishable until t , inclusive.

Definition 6.25 (Retainable Subsets). Given a context γ , a protocol $P \in \mathbb{P}_\gamma$, a run $r \in R_\gamma(P)$ and a time $t \in \mathbb{T}$, we define the “ t -retainable” subsets of $ND_\gamma(r)$ as

$$RND_\gamma^P(r, t) \triangleq \{ND_\gamma(r', t) \mid r' \stackrel{t}{\subseteq} r\} \subseteq 2^{ND_\gamma(r, t)}.$$

Furthermore, for every $E \in RND_\gamma^P(r, t)$, we denote

$$r \stackrel{t}{\cap} E \triangleq \{r' \stackrel{t}{\subseteq} r \mid ND_\gamma(r', t) = E\} \neq \emptyset,$$

(Again, we omit P and γ from this notation as they will be inferred from the discussion.) We sometimes slightly abuse notation by using $r \stackrel{t}{\cap} E$ to refer to one such run and not to the whole set, if the choice of representative is inconsequential. (This is often the case, as $r \stackrel{t}{\cap} E$ is an equivalence class of $\stackrel{t}{\subseteq}$.)

Remark 6.26. Let γ be a context, let $P \in \mathbb{P}_\gamma$, let $r \in R_\gamma(P)$ and let $t \in \mathbb{T}$. By the above definitions:

- $\forall t \in \mathbb{T} : ND_\gamma(r, t) \in RND_\gamma^P(r, t)$, and $r \in r \stackrel{t}{\cap} ND_\gamma(r, t)$.
- If $E \in RND_\gamma^P(r, t)$ and if $E' \in RND_\gamma^P(r \stackrel{t}{\cap} E, t')$ for some $t' \in \mathbb{T}$ s.t. $t' \leq t$, then by definition, $E' \subseteq E$, $E' \in RND_\gamma^P(r, t')$ and $(r \stackrel{t}{\cap} E) \stackrel{t'}{\cap} E' = r \stackrel{t'}{\cap} E'$.

Claim 6.27. *Let γ be a protocol, let $P \in \mathbb{P}_\gamma$, let $r \in R_\gamma(P)$ and let $t, t' \in \mathbb{T}$. If $t' \leq t$, then $RND_\gamma^P(r, t') \subseteq RND_\gamma^P(r, t)$. Furthermore, for every $E \in RND_\gamma^P(r, t')$, there exists $r' \in r \overset{t}{\cap} E$ s.t. $ND_\gamma(r') \cap \tilde{E}_\gamma \subseteq E$.*

Proof Sketch. Let $E \in RND_\gamma^P(r, t')$. For the continuous-time model presented in Appendix A, the claim follows by applying the “no foresight” property to $r \overset{t'}{\cap} E$ at t' and with $d = t$. For the discrete-time model presented in Chapter 3, we construct a run $r' \in r \overset{t}{\cap} E$ s.t. $ND_\gamma(r') \cap \tilde{E}_\gamma \subseteq E$, as follows: r' is identical to $r \overset{t'}{\cap} E$ until t' , inclusive. After t' , the agents behave in r' according to P , and the environment triggers no more ND events, except for deliveries of sent messages that have an infinite bound guarantee (as such non-deterministic deliveries must be triggered at some time during the run, for the run to be legal). Each such message is delivered at $\max\{t+1, t''+1\}$, where t'' is the sending time of this message. It is straightforward to inductively check that the resulting run r' is well defined and legal — we omit this cumbersome check, which runs along similar lines of some of the proofs from [6], in favor of the many, more interesting, pages ahead. \square

We now present and adapt some machinery developed by Ben-Zvi and Moses in their analysis[6, 5] of the ordered response problem. Their discrete-time analysis essentially shows the following lemma, which we rephrase using our notation. For the continuous-time model presented in Appendix A, the first part of this lemma is equivalent to the “no extrasensory perception” property, and its second part follows from the definition of a run.

Lemma 6.28. *Let γ be a context, let $P \in \mathbb{P}_\gamma$ and let $r \in R_\gamma(P)$. For every $t \in \mathbb{T}$,*

$$RND_\gamma^P(r, t) \supseteq \{PND_\gamma^r(i, t) \mid i \in \mathbb{I}_\gamma\},$$

and for each $i \in \mathbb{I}_\gamma$, the state of i at t is identical in r and in all the runs $r \overset{t}{\cap} PND_\gamma^r(i, t)$.

By applying Claim 6.27, we obtain the following generalization of Lemma 6.28.

Corollary 6.29. *Let γ be a context, let $P \in \mathbb{P}_\gamma$ and let $r \in R_\gamma(P)$. For every $t \in \mathbb{T}$,*

$$RND_\gamma^P(r, t) \supseteq \{PND_\gamma^r(i, t') \mid i \in \mathbb{I}_\gamma \ \& \ t' \leq t\},$$

and for each $i \in \mathbb{I}_\gamma$ and each $t' \leq t$, the state of i at t' is identical in r and in all the runs $r \overset{t}{\cap} PND_\gamma^r(i, t')$.

We note, without a proof, that this result can be further generalized as follows, at least for the cases listed below.

Corollary 6.30. *In a discrete-time model, or in a continuous-time model with finitely many agents, let γ be a context, let $P \in \mathbb{P}_\gamma$ and let $r \in R_\gamma(P)$.^{*} For every $t \in \mathbb{T}$,*

$$RND_\gamma^P(r, t) \supseteq \{E \subseteq ND_\gamma(r, t) \mid \forall e \in E : E \supseteq PND_\gamma^r(e)\},$$

with equality if P is a full-information protocol. Furthermore, for every $E \in RND_\gamma^P(r, t)$ and for every $(i, t') \in \mathbb{I}_\gamma \times \mathbb{T}$, if $PND_\gamma^r(i, t') \subseteq E$, then the state of i at t' is identical in r and in all the runs $r \overset{t}{\cap} E$. (If P is a full-information protocol, then the converse holds as well.)

We do not require Corollary 6.30, though, as Corollary 6.29 suffices for all the proofs that we give below.

6.3 Analyzing Timely-Coordinated Response

We now turn to define the structure that stands at the heart of our syncausal analysis of the timely-coordinated response problem.

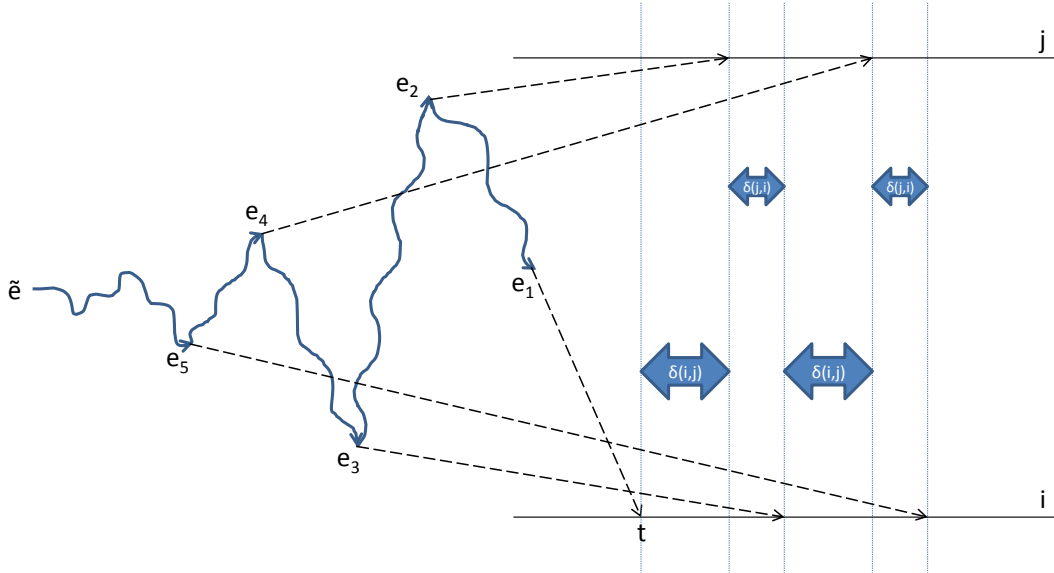


Figure 4: $(e_m)_{m=1}^5$ is an $((i, j, i, j, i), \delta)$ -traversing \tilde{e} -centipede by t .

^{*} As discussed in Appendix A, for certain “nice” protocols P , the requirement for only finitely many agents in a continuous-time model may be relaxed to the requirement that $\inf(b_\gamma) > 0$.

Definition 6.31 (Path-Traversing Centipede — see Figure 4). *Given a TCR-spec $(\gamma, \tilde{e}, I, \delta)$, a path $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$ and a run $r \in \mathcal{R}_\gamma$, we call an n -tuple of ND events $\bar{e} \in ND_\gamma(r)^n$ a “ (\bar{p}, δ) -traversing \tilde{e} -centipede” by t if the following hold:*

- $\tilde{e} \xrightarrow[\gamma]{r} e_n$ and $\forall m \in [n-1] : e_{m+1} \xrightarrow[\gamma]{r} e_m$.
- $\forall m \in [n] : e_m \xrightarrow[\gamma]{} (p_m, t + L_{G_\delta}((p_k)_{k=1}^m))$. We call $\{(p_m, t + L_{G_\delta}((p_k)_{k=1}^m))\}_{m=1}^n$ the set of “end nodes” of this centipede.

Remark 6.32. $\bar{e}^{rev} = (e_{n-m+1})_{m=1}^n$ is an \tilde{e} -centipede for \bar{p}^{rev} (as a tuple of agents) by \bar{t}^{rev} , where for every $m \in [n]$, $t_m \triangleq t + L_{G_\delta}((p_k)_{k=1}^m)$. Thus, $t_{m+1} = t_m + \delta(p_m, p_{m+1})$ for every $m \in [n-1]$.

Remark 6.33. Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec, let $\bar{p} \in \mathcal{P}(G_\delta)$ and let $r \in \mathcal{R}_\gamma$. By the above definition:

- No (\bar{p}, δ) -traversing \tilde{e} -centipede exists in r , if \bar{p} traverses an edge with a weight of $-\infty$ in G_δ .
- Any (\bar{p}, δ) -traversing \tilde{e} -centipede by t in r is also a (\bar{p}, δ') -traversing \tilde{e} -centipede by t' in r , for every $\delta' \geq \delta$ and every $t' \geq t$, by the locality property of bound guarantee. This justifies the phrasing “path-traversing centipede by t ”.
- Let $\bar{e} = (e_m)_{m=1}^n$ be a (\bar{p}, δ) -traversing \tilde{e} -centipede by t in r , then $(e_m)_{m=k}^n$ is a $((p_m)_{m=k}^n, \delta)$ -traversing \tilde{e} -centipede by $t + L_{G_\delta}((p_m)_{m=1}^k)$ in r , for every $k \in [n]$. We call this path-traversing centipede the “ k -suffix” of \bar{e} .

The following theorem, once stated, may be proven using the tools that are applied in [8] for proving Theorem 6.21. We provide a somewhat different and more concise proof here, also for the sake of emphasizing the fact that the approach studied in this chapter requires no direct use of the concept of knowledge.

Theorem 6.34 (Path-Traversing Centipede). *Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec and let $P \in TCR_\gamma\langle \tilde{e}, I, \delta \rangle$. Each $r \in R_\gamma^\tilde{e}(P)$ contains a (\bar{p}, δ) -traversing \tilde{e} -centipede by $t_r(p_1)$, for every $\bar{p} \in \mathcal{P}(G_\delta)$.*

Proof. By induction on n , the number of vertices in \bar{p} . ($\bar{p} = (p_m)_{m=1}^n$.)

Base: If $n = 1$, denote $i \triangleq p_1$ (and thus, $\bar{p} = (i)$). Since $r \in R_\gamma^\tilde{e}(P)$, we claim that $\tilde{e} \xrightarrow[\gamma]{r} (i, t_r(i))$. Indeed, by Corollary 6.29 and by Claim 6.27, there exists a run $r' \in R_\gamma(P)$ for which $ND_\gamma(r', t_r(i)) = PND_\gamma^r(i, t_r(i))$ and in which the only

occurring external inputs are those that are in $PND_\gamma^r(i, t_r(i))$. Furthermore, both the state of i , and the events observed by it, are identical in r and in r' up to and including $t_r(i)$, and thus $t_{r'}(i) = t_r(i) < \infty$. By correctness of P , this implies $r' \in R_\gamma^{\tilde{e}}(P)$, and thus $\tilde{e} \in ND_\gamma(r) \cap \tilde{E}_\gamma \subseteq PND_\gamma^r(i, t_r(i))$, as required. Thus, there exists a syncausal path in r from \tilde{e} to $(i, t_r(i))$. Denote the latest among the ND event along this path by $e \in ND_\gamma(r)$. By definitions of syncausality and of bound guarantee, $\tilde{e} \xrightarrow[r]{\gamma} e \xrightarrow[r]{\gamma} (i, t_r(i))$. Thus, (e) is a path-traversing centipede as required.

Induction step: Assume $n \geq 2$. Set $i \triangleq p_1$, $j \triangleq p_2$, $E \triangleq PND_\gamma^r(i, t_r(i))$ and $t \triangleq \max\{t_r(i), t_r(i) + \delta(i, j)\} < \infty$. As $t_r(i) \leq t$, by Corollary 6.29 we obtain that $E \in RND_\gamma^P(r, t)$, and that the state of i at $t_r(i)$ is the same in $r \overset{t}{\cap} E$ and in r . Therefore, $t_{r \overset{t}{\cap} E}(i) = t_r(i)$, and thus, by correctness of P , we obtain

$$t_{r \overset{t}{\cap} E}(j) \leq t_{r \overset{t}{\cap} E}(i) + \delta(i, j) = t_r(i) + \delta(i, j).$$

By the induction hypothesis, there exists a $((p_m)_{m=2}^n, \delta)$ -traversing \tilde{e} -centipede $(e_m)_{m=2}^n$ by $t_{r \overset{t}{\cap} E}(j)$ (and thus, by Remark 6.33, by $t_r(i) + \delta(i, j)$) in $r \overset{t}{\cap} E$, and thus also in r (as $t_r(i) + \delta(i, j) \leq t$). To complete our proof, we note that

$$\begin{aligned} e_2 \in PND_\gamma^{r \overset{t}{\cap} E}(j, t_r(i) + \delta(i, j)) &\subseteq && \text{by definition of } PND \\ &\subseteq ND_\gamma(r \overset{t}{\cap} E, t_r(i) + \delta(i, j)) &\subseteq & \text{as } t_r(i) + \delta(i, j) \leq t \\ &\subseteq ND_\gamma(r \overset{t}{\cap} E, t) = E = PND_\gamma^r(i, t_r(i)). \end{aligned}$$

and therefore $e_2 \xrightarrow[r]{\gamma} (i, t_r(i))$. As in the induction base, there exists $e_1 \in ND_\gamma(r)$ s.t. $e_2 \xrightarrow[r]{\gamma} e_1 \xrightarrow[r]{\gamma} (i, t_r(i))$. Thus, $(e_m)_{m=1}^n$ is a path-traversing centipede as required. \square

It should be noted that by Corollary 5.14, a $(\bar{p}, \hat{\delta})$ -traversing \tilde{e} -centipede is also implied by Theorem 6.34 for every $\bar{p} \in \mathcal{P}(G_{\hat{\delta}})$ under the conditions of that theorem. Furthermore, this result is at least as strong as the verbatim result of that theorem for δ , by minimality of the canonical form, by Remark 6.33 and as $\mathcal{P}(G_\delta) \subseteq \mathcal{P}(G_{\hat{\delta}})$. It may be readily verified that if the distance between every pair of agents in G_δ is attained, then these results are in fact equivalent, as any path-traversing centipede guaranteed by Theorem 6.34 for $\hat{\delta}$ is a (possibly trivial) subcentipede (i.e. subtuple) of a path-traversing centipede directly guaranteed by it for δ (for a possibly different path). Henceforth, whenever minimizing the times of the end nodes of the guaranteed path-traversing centipede is of the essence (as is the case in e.g. Corollaries 9.4

and 9.5 and Claim 9.14), we indeed apply Theorem 6.34 using $\hat{\delta}$.

We now apply Theorem 6.34 to deduce an optimal response logic for the timely-coordinated response problem in shared-clock models.

Corollary 6.35. *In a shared-clock model, let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec. An optimal response logic for solving $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is, for every $i \in I$: “respond at the earliest time by which a (\bar{p}, δ) -traversing \tilde{e} -centipede exists for every path $\bar{p} \in \mathcal{P}(G_\delta)$ starting at $p_1 = i$ ”.*

Proof. Assume that $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable and let $P \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle$. W.l.o.g., P is a full-information protocol. Let P' be the (full-information) protocol obtained by endowing P with the above-defined response logic. We first prove the optimality of P' in each triggered run, and then prove that it indeed solves $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$.

Let $r \in R_\gamma^\tilde{e}(P)$ and $r' \in R_\gamma^\tilde{e}(P')$ be two runs matched under the natural isomorphism between $R_\gamma(P)$ and $R_\gamma(P')$ and let $i \in I$. By Theorem 6.34, all path-traversing centipedes required for i to respond according to P' exist in r (and hence in r') by $t_r(i)$ (at the latest), and therefore $t_{r'}(i) \leq t_r(i)$. (The fact that P' is a full-information protocol, together with the existence of a shared clock, guarantees that if such path-traversing centipedes exist by some $t \in \mathbb{T}$, then i can deduce this at t . For the continuous-time model presented in Appendix A, the fact that i is enabled at every supremum of times at which it observes events allows i to respond at the required time.)

We now prove that $P' \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle$: Let $r' \in R_\gamma(P')$. Obviously, no path-traversing \tilde{e} -centipedes exist in r' if \tilde{e} does not occur. Therefore, if $r' \notin R_\gamma^\tilde{e}(P')$ then $t_{r'} \equiv \infty$. We are left with the case in which $r' \in R_\gamma^\tilde{e}(P')$. Denote by $r \in R_\gamma^\tilde{e}(P)$ the run of P matching r' under the natural isomorphism between $R_\gamma(P)$ and $R_\gamma(P')$. By the first part of this proof, $t_{r'} \leq t_r < \infty$. Let $(i, j) \in I^2$ s.t. $\delta(i, j) < \infty$. For every $\bar{p} \in \mathcal{P}(G_\delta)$ s.t. $p_1 = j$, denote by $\bar{p}' \in \mathcal{P}(G_\delta)$ the path commencing at i and whose 2-suffix is \bar{p} . By definition of P' , a (\bar{p}', δ) -traversing \tilde{e} -centipede exists in r' by $t_{r'}(i)$ (at the latest), and thus its 2-suffix, which is a (\bar{p}, δ) -traversing \tilde{e} -centipede, exists in r' by $t_{r'}(i) + \delta(i, j)$ (at the latest). Hence, all path-traversing centipedes required for j 's response according to P' exist in r' by $t_{r'}(i) + \delta(i, j)$, and thus $t_{r'}(j) \leq t_{r'}(i) + \delta(i, j)$, as required.* \square

We conclude this chapter with a practical note that motivates some of our dis-

* The attentive reader may notice a conceptual similarity between this argument and the second part of the proof of Lemma 5.7.

cussion in Chapter 9. In that chapter, we derive somewhat more practical results from the above discussion, for some naturally-occurring models that we define.

The results of Ben-Zvi and Moses that are surveyed in the beginning of this chapter imply that for each of the coordinated response problems they have studied, it is enough for an agent i to deduce the existence of a single, simple, syncausal structure in order to respond according to the optimal response logic for this problem. In contrast, from the definition of the optimal response logic for the timely-coordinated response problem from Corollary 6.35, it may seem that in the case of a general constraining function δ (i.e. δ that does not reduce to e.g. one of the special cases studied by Ben-Zvi and Moses), for an agent i to respond according to this logic, i is always required to check for infinitely many, arbitrarily long, path-traversing centipedes (using infinitely many facts stored in the memory/state of i). While in the general case this is true, the following remark shows that for any constraining function δ , in some cases finitely many syncausal structures may imply the existence of all (infinitely many) path-traversing centipedes required for i 's response. Conceptually, this means that finitely many checks (of finitely many facts) may provide enough information for i to respond at a specific time according to this logic. In Chapter 9, we show that under certain practical assumptions, finitely many not-much-more-complicated checks always suffice.

Remark 6.36. *Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec s.t. δ is implementable, and let $r \in \mathcal{R}_\gamma$. If $e \in ND_\gamma(r)$ is an \tilde{e} -broom for I in r by $\bar{t} \in \mathbb{T}^I$ s.t. $\sup(\bar{t}) < \infty$, then for every $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$, $(e)^n$ is a (\bar{p}, δ) -traversing \tilde{e} -centipede in r by*

$$\max_{k \in [n]} \{t_{p_k} - L_{G_\delta}((p_m)_{m=1}^k)\} \leq \max_{k \in [n]} \{t_{p_k} - \hat{\delta}(p_1, p_k)\} \leq \sup(\bar{t}) - \inf(\hat{\delta}|_{\{p_1\} \times I}) < \infty.$$

*In particular, for every agent $i \in I$, $(e)^n$ is a (\bar{p}, δ) -traversing \tilde{e} -centipede by $\sup(\bar{t}) - \inf(\hat{\delta}|_{\{i\} \times I})$ (which is finite by implementability of δ and by Lemma 5.7), for every $n \in \mathbb{N}$ and every $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$ starting at $p_1 = i$. Thus, by this time i will have received information guaranteeing that e had occurred, which will have given i enough information in order to respond by that time according to the optimal response logic presented in Corollary 6.35.**

* We note that it is possible to construct an alternative argument as to why there exists, under certain conditions, a solving finite-memory protocol according to which each i responds by that time. Such an argument may be constructed by combining a variation of the second part of Theorem 6.12, with the second part of Lemma 5.7 and with the second part of Claim 5.1. Conversely, Remark 6.36 may be used to construct an alternative proof for parts of Lemma 5.7.

Chapter 7

The Fixed-Point Approach

We now set aside, for the moment, the results of Chapter 6 and embark on a parallel, independent analysis of the timely-coordinated response problem using fixed-point analysis. While the basis of this analysis follows the lines of [11, Section 11.6], we formalize it here using events, along the approach of Aumann[3], instead of using the temporal-epistemic logic tools used in [11]. (The treatment in either form is analogous, although it is somewhat more concise for our case with the notation used below, which facilitates the study of fixed points.)

7.1 Background

In this section, we survey previous definitions and results from [11], upon which our analysis below is founded, reformulating them using events, and adapting them to our notation.

7.1.1 Events, Knowledge and Common Knowledge

In order to begin our discussion, we define the space in which we work.

Definition 7.1 (Space). *Let γ be a context and let $R \subseteq \mathcal{R}_\gamma$. We define $\Omega_R \triangleq R \times \mathbb{T}$ and $\mathcal{F}_R \triangleq 2^{\Omega_R}$.*

As in probability theory, we represent events using the set of points (i.e. run-time pairs) in which they hold. (In contrast to probability theory, though, we do not need to define a measure on the set of events, so we choose to allow any subset of Ω_R to constitute an event.) For example, we may define the event “ i is responding” for

some $i \in \mathbb{I}_\gamma$, which is formally associated with all points $(r, t) \in \Omega_R$ s.t. i responds at t in r .

We now incorporate the concept of knowledge into our discussion. Given an event $\psi \in \mathcal{F}_R$, we wish to define, for some agent i , the event “ i knows that ψ holds” (e.g. i knows that j is responding). Choosing how to formalize a concept as abstract and as subjective as knowledge is not a simple issue. We present below what has become a standard definition for knowledge, and avoid discussing its relation to the abstract, philosophical, concept of knowledge. Intuitively, by this definition, at (r, t) i knows ψ iff ψ holds at all possible points (s, t') that i cannot distinguish from (r, t) .

Definition 7.2 (Knowledge). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $i \in \mathbb{I}_\gamma$.*

1. *We partition Ω_R into equivalence classes according to the state of i , s.t. $p, q \in \Omega_R$ are in the same equivalence class iff the state of i is the same in p and in q . In a shared-clock model, we additionally demand that the time be the same at p and at q . For $p \in \Omega_R$, we denote the equivalence class of p by $\mathcal{S}_i(p)$.*

2. *Define*

$$K_i : \mathcal{F}_R \rightarrow \mathcal{F}_R$$

$$\psi \mapsto \{p \in \Omega_R \mid \mathcal{S}_i(p) \subseteq \psi\}.$$

While both definitions, of \mathcal{S}_i and of K_i , depend on R , we omit R from these notations, for readability, as the set of runs will be clear from the discussion. We follow this convention when presenting some other definitions in this, and in the following, chapter as well.

We now present a few immediate (and well-known) properties of the knowledge operator. The first one, sometimes referred to as the “Truth Axiom for Knowledge”, intuitively means that “whenever anyone knows something, then it is true”. The second property, sometimes referred to as the “Positive Introspection Axiom”, which intuitively means “whenever i knows something, then i knows that it knows it”, has been the subject of quite a few philosophical discussions. As we have done when presenting the definition for knowledge, and as we will continue to do below, we present it, and avoid discussing any philosophical consequences thereof.

Remark 7.3. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $i \in \mathbb{I}_\gamma$. By Definition 7.2, the knowledge operator K_i satisfies:*

- *Truth Axiom for Knowledge:* $K_i(\psi) \subseteq \psi$, for every event $\psi \in \mathcal{F}_R$.
- *Positive Introspection Axiom:* $K_i(K_i(\psi)) = K_i(\psi)$, for every event $\psi \in \mathcal{F}_R$.
- *Monotonicity:* $\psi \subseteq \phi \Rightarrow K_i(\psi) \subseteq K_i(\phi)$, for every two events $\psi, \phi \in \mathcal{F}_R$.
- *K_i commutes with intersection:* $K_i(\cap \Psi) = \cap \{K_i(\psi) \mid \psi \in \Psi\}$, for every set of events $\Psi \subseteq \mathcal{F}_R$.

We now build upon Definition 7.2 and define the notion of “everybody knows”.

Definition 7.4. Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$ be a set of agents. Define

$$E_I : \mathcal{F}_R \rightarrow \mathcal{F}_R$$

$$\psi \mapsto \bigcap_{i \in I} K_i(\psi).$$

A truth axiom, analogous to the one presented in Remark 7.3, readily holds for E_I as well. In addition, E_I is monotone and commutes with intersection. However, it is not idempotent.

We are now ready to define common knowledge. One classic, constructive definition of common knowledge[12] is the following, defining that an event is common knowledge to a set of agents when all know it, all know that all know it, etc.

Definition 7.5 (Common Knowledge). Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. Define

$$C_I : \mathcal{F}_R \rightarrow \mathcal{F}_R$$

$$\psi \mapsto \bigcap_{n=1}^{\infty} E_I^n(\psi),$$

where $E_I^0(\psi) = \psi$ and $E_I^n(\psi) = E_I(E_I^{n-1}(\psi))$ for every $n \in \mathbb{N}$.

It may be readily verified that the common knowledge operator satisfies the obvious analogues of all properties of the knowledge operator that are presented in Remark 7.3 (including idempotence*).

A classic result[14] regarding common knowledge is that it relates tightly to simultaneous response, in the sense that e.g. in order to coordinate a simultaneous response among a set of agents, they must all have common knowledge of the

* In fact, C_I is the (coordinate-wise) greatest idempotent operator s.t. $C_I \subseteq E_I$.

response when it occurs. Conversely, whenever common knowledge of a fact arises among a set of agents, it does so simultaneously for all agents. Definition 7.9 and Theorem 7.10 below formalize this intuition, but before we present them, we turn to a few more definitions.

As noted above, the truth axiom for knowledge, presented in Remark 7.3, implies that whenever some fact is known to someone, the fact is true as well. Certain events, such as, for some agent i , “ i is responding (right now)”, have the converse property as well, i.e. they are known to i whenever they hold. The following definition characterises such events.

Definition 7.6 (Local Event). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $i \in \mathbb{I}_\gamma$. An event $\psi \in \mathcal{F}_R$ is said to be “local” to i if $K_i(\psi) = \psi$.*

Remark 7.7. *By the positive introspection axiom presented in Remark 7.3, $K_i(\psi)$ is local to i , for every $i \in \mathbb{I}_\gamma$ and for every $\psi \in \mathcal{F}_R$.*

An important property of events that are local to i , for some agent $i \in \mathbb{I}_\gamma$, is that i may act upon them, i.e. the response logic of i in a protocol may be defined by specifying that i should respond whenever some given local event for i holds. (i may do that, as locality of this event guarantees that whether it holds or not at some time $t \in \mathbb{T}$ is determined by the state of i at t .) As we are interested in coordination, though, we are usually interested in specifying a joint response logic for a set of agents $I \subseteq \mathbb{I}_\gamma$, i.e. a response logic for each $i \in I$, in which the response times of the various agents are coordinated in some way. One way to specify such a joint response logic, therefore, is to specify, for each $i \in I$, a local event for i . Such a collection of specifications is called an “ensemble”.

Definition 7.8 (Event Ensemble). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. An I -tuple of events $\bar{\mathbf{e}} = (\mathbf{e}_i)_{i \in I} \in \mathcal{F}_R^I$ is called an ensemble if \mathbf{e}_i is local to i for each $i \in I$.*

It should be noted, though, that ensembles are useful beyond specifying response logics, as they may be used to study the coordination of passive events as well, i.e. coordination of times at which agents become aware of some fact, or at which they observe an event (e.g. receive some message).

We are now ready to survey the results of [14] relating common knowledge and simultaneity, as formulated for ensembles in [11, Section 11.6]. While phrasing the following theorem, and henceforth, we use the following shorthand notation: $\cup \bar{\xi} \triangleq \cup_{i \in I} \xi_i$, for every $\bar{\xi} = \xi_{i \in I} \in \mathcal{F}_R^I$.

Definition 7.9 (Perfect Coordination). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. An ensemble $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is said to be “perfectly coordinated” if $\mathbf{e}_i = \mathbf{e}_j$ for every $i, j \in I$.*

Theorem 7.10. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$.*

1. *For every event ψ , the ensemble $(K_i(C_I(\psi)))_{i \in I}$ is perfectly coordinated.*
2. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is a perfectly coordinated ensemble, then $\mathbf{e}_i \subseteq K_i(C_I(\cup \bar{\mathbf{e}}))$ for every $i \in I$.*
3. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is a perfectly coordinated ensemble, then $\cup \bar{\mathbf{e}} \subseteq C_I(\cup \bar{\mathbf{e}})$.*

In Chapter 8, we show that the analysis of Chapter 6 has led us, in a sense, to a definition that is similar to Definition 7.5.

7.1.2 Fixed-Point Analysis

Another classic definition[19] for common knowledge, which is known to be equivalent, is the following, defining it as the greatest fixed point of a function on \mathcal{F}_R .

Theorem 7.11 (Common Knowledge as a Greatest Fixed Point). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. $C_I(\psi)$ is the greatest fixed point of the function $x \mapsto E_I(\psi \cap x)$, for every event $\psi \in \mathcal{F}_R$.*

Based on this definition, Moses and Halpern[14] defined two variants of Common Knowledge, matching two weaker forms of coordination. We now present their results, as formulated for ensembles in [11, Section 11.6].

Definition 7.12 (Eventual Coordination). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. An ensemble $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is said to be “eventually coordinated” if for every $i, j \in I$ and for every $(r, t) \in \mathbf{e}_i$, there exists $t' \in \mathbb{T}$ s.t. $(r, t') \in \mathbf{e}_j$.*

Definition 7.13. *Let γ be a context and let $R \subseteq \mathcal{R}_\gamma$. Define*

$$\begin{aligned} \diamond : \mathcal{F}_R &\rightarrow \mathcal{F}_R \\ \psi &\mapsto \{(r, t) \in \Omega_R \mid \exists t' \in \mathbb{T} : (r, t') \in \psi\}^* \end{aligned}$$

($\diamond(\psi)$ is the event “ ψ eventually holds at some time during the current run, be it past, present or future.”)

* We use the symbol \diamond instead of the standard temporal logic notation \Diamond , in order to emphasize that t' may be smaller than t .

Theorem 7.14. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$.*

1. *For every $\psi \in \mathcal{F}_R$, the function $x \mapsto \bigcap_{i \in I} \Diamond (K_i(\psi \cap x))$ has a greatest fixed point. Denote it by $C_I^\Diamond(\psi)$ (“eventual common knowledge” of ψ by I).*
2. *For every event ψ , the ensemble $(K_i(C_I^\Diamond(\psi)))_{i \in I}$ is eventually coordinated.*
3. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is an eventually coordinated ensemble, then $\mathbf{e}_i \subseteq K_i(C_I^\Diamond(\cup \bar{\mathbf{e}}))$ for every $i \in I$.*
4. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is an eventually coordinated ensemble, then $\cup \bar{\mathbf{e}} \subseteq C_I^\Diamond(\cup \bar{\mathbf{e}})$.*

Another variant of common knowledge, also defined and studied by Halpern and Moses, relates to an approximation of perfect coordination.

Definition 7.15 (ε -Coordination). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$, let $I \subseteq \mathbb{I}_\gamma$ and let $\varepsilon \geq 0$. An ensemble $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is said to be “ ε -coordinated” if for every $i \in I$ and for every $(r, t) \in \mathbf{e}_i$, there exists an interval $T \subseteq \mathbb{T}$ of length at most ε , s.t. $t \in T$ and s.t. for every $j \in I$ there exists $t' \in T$ s.t. $(r, t') \in \mathbf{e}_j$.*

We note that 0-coordination is the same as perfect coordination, and thus the following theorem also implies Theorem 7.10 as a special case thereof.

Theorem 7.16. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$, let $\varepsilon \geq 0$ and let $I \subseteq \mathbb{I}_\gamma$. Define*

$$E_I^\varepsilon : \mathcal{F}_R \rightarrow \mathcal{F}_R$$

$$\psi \mapsto \left\{ (r, t) \in \Omega_R \mid \exists T \subseteq \mathbb{T} : \begin{array}{l} t \in T \ \& \ \sup\{T - T\} \leq \varepsilon \ \& \\ \forall i \in I \ \exists t' \in T : (r, t') \in K_i(\psi) \end{array} \right\}.$$

1. *For every $\psi \in \mathcal{F}_R$, the function $x \mapsto E_I^\varepsilon(\psi \cap x)$ has a greatest fixed point. Denote it by $C_I^\varepsilon(\psi)$ (“ ε -common knowledge” of ψ by I).*
2. *For every event ψ , the ensemble $(K_i(C_I^\varepsilon(\psi)))_{i \in I}$ is ε -coordinated.*
3. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is an ε -coordinated ensemble, then $\mathbf{e}_i \subseteq K_i(C_I^\varepsilon(\cup \bar{\mathbf{e}}))$ for every $i \in I$.*
4. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is an ε -coordinated ensemble, then $\cup \bar{\mathbf{e}} \subseteq C_I^\varepsilon(\cup \bar{\mathbf{e}})$.*

The attentive reader may notice, by now, a pattern forming in the similarity between Definition 7.9 and the function defined in Theorem 7.11, between Definition 7.12 and the function defined in Theorem 7.14 and between Definition 7.15 and the function defined in Theorem 7.16.

7.2 δ -Common Knowledge

Having completed our survey of some previous, relevant, results and definitions, we are now ready to start extending them. Recall that the interdependencies between the response times of different agents in the timely-coordinated response problem are captured by an implementation-spec (I, δ) , i.e. a set of agents I and a constraining function $\delta : I^2 \rightarrow \Delta$ from ordered pairs of distinct agents, to maximum allowed time differences. Before we turn to analyze the timely-coordinated response problem in the next section, we first define a form of coordination exhibiting similar constraints, and analyze it.

Definition 7.17 (δ -Coordination). *We call a quadruplet (γ, R, I, δ) a “ δ -coordination-spec”*, if γ is a context, $R \subseteq \mathcal{R}_\gamma$ is a set of runs, and (I, δ) is an implementation-spec s.t. $I \subseteq \mathbb{I}_\gamma$. Given a δ -coordination-spec (γ, R, I, δ) , we say that an ensemble $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is “ δ -coordinated” if for every $(i, j) \in I^2$ and for every $(r, t) \in \mathbf{e}_i$, there exists $t' \in T$ s.t. $t' \leq t + \delta(i, j)$ and $(r, t') \in \mathbf{e}_j$.*

It should be noted that from this point on, whenever dealing with coordinated ensembles, we always assume, for ease of presentation, that $|I| > 1$, i.e. that the ensemble is defined over more than a single agent. Adjusting our results for the case in which this does not hold is neither hard, nor interesting.

Before we commence our analysis of δ -coordination, we define, given an event $\psi \in \mathcal{F}_R$, notation standing for the event “ ψ holds at some (past, present, or future) time, no later than ε time units from now”.

Definition 7.18. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $\varepsilon \in \Delta$. We define*

$$\begin{aligned} \odot^{\leq \varepsilon} : \mathcal{F}_R &\rightarrow \mathcal{F}_R \\ \psi &\mapsto \{(r, t) \in \Omega_R \mid \exists t' \subseteq \mathbb{T} : t' \leq t + \varepsilon \ \& \ (r, t') \in \psi\}^\S \end{aligned}$$

Remark 7.19. *By Definition 7.18:*

$$\bullet \odot^{\leq \infty} = \diamond.$$

* Note the difference between the italic δ that indicates a specific constraining function, and the roman (i.e. upright) δ that generally refers to the form of coordination that we define, regardless of any concrete constraining function. For example, δ -coordination, for a specific constraining function δ , is an instance of δ -coordination.

[§] As with our usage of \diamond , we use the symbol \odot instead of the standard temporal logic notation \circ , in order to emphasize that t' may be smaller than t .

- $\odot^{\leq -\infty}(\psi) = \emptyset$, for every $\psi \in \mathcal{F}_R$.
- $\odot^{\leq 0}(\psi)$, for an event $\psi \in \mathcal{F}_R$, means “ ψ has occurred, either now or in the past”.
- *Additivity:* $\odot^{\leq \varepsilon_1}(\odot^{\leq \varepsilon_2}(\psi)) = \odot^{\leq \varepsilon_1 + \varepsilon_2}(\psi)$ for every $\varepsilon_1, \varepsilon_2 \in \Delta \setminus \{-\infty\}$ and for every event $\psi \in \mathcal{F}_R$.
- *Monotonicity:* $(\varepsilon_1 \leq \varepsilon_2 \ \& \ \psi \subseteq \phi) \Rightarrow \odot^{\leq \varepsilon_1}(\psi) \subseteq \odot^{\leq \varepsilon_2}(\phi)$, for every $\varepsilon_1, \varepsilon_2 \in \Delta$ and for every two events $\psi, \phi \in \mathcal{F}_R$.
- $\odot^{\leq \varepsilon}(\cap \Psi) \subseteq \cap \{\odot^{\leq \varepsilon}(\psi) \mid \psi \in \Psi\}$, for every $\varepsilon \in \Delta$ and for every set of events $\Psi \subseteq \mathcal{F}_R$.

We are now ready to analyze δ -coordination along the lines of the results surveyed in the previous section. First, we define a common-knowledge analogue for this case. As in the results presented in the previous section, given an event $\psi \in \mathcal{F}_R$, we use ψ to define a function f_ψ^δ , s.t. knowledge of the greatest fixed point of f_ψ^δ by each agent constitutes a δ -coordinated ensemble with several desired properties. Nonetheless, we face several additional technical challenges along the way. A main technical challenge is that δ -coordination lacks the symmetry among the different agents in I , which manifests in the common knowledge variants presented in the previous section. Thus, in general there is no natural way to define a single event ψ for which $(K_i(\psi))_{i \in I}$ is the ensemble we are looking for, i.e. our ensemble should be defined in terms of knowledge of different events for different agents. For this reason, we somewhat generalize our strategy: Instead of searching for a fixed point of a function on \mathcal{F}_R , we define a function on \mathcal{F}_R^I — the set of I -tuples of events. We denote the greatest fixed point of this function by $C_I^\delta(\psi)$ (this is an I -tuple of events), and show that $(K_i(C_I^\delta(\psi)_i))_{i \in I}$ is the desired ensemble, i.e. each coordinate of this fixed point is the event that $i \in I$ should know in this ensemble.* To our knowledge, such a technique was never utilized in this field before.

Before we define the above-described function, we define a lattice structure on \mathcal{F}_R^I , which gives precise meaning to the concept of a greatest fixed point of a function on \mathcal{F}_R^I .

* While vectorial fixed points may alternatively be captured by nested fixed points [2, Chapter 1], in our case we argue that the vectorial representation better parallels the underlying intuition.

Definition 7.20 (Lattice Structure on \mathcal{F}_R^I). *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $I \subseteq \mathbb{I}_\gamma$. We define the following lattice structure on \mathcal{F}_R^I : (In the following definitions, $\bar{\varphi} \triangleq (\varphi_i)_{i \in I} \in \mathcal{F}_R^I$ and $\bar{\xi} \triangleq (\xi_i)_{i \in I} \in \mathcal{F}_R^I$.)*

1. *Order:* $\bar{\varphi} \leq \bar{\xi} \iff \forall i \in I : \varphi_i \subseteq \xi_i$.

2. *Join:* $\bar{\varphi} \vee \bar{\xi} \triangleq (\varphi_i \cup \xi_i)_{i \in I}$.

3. *Meet:* $\bar{\varphi} \wedge \bar{\xi} \triangleq (\varphi_i \cap \xi_i)_{i \in I}$.

Remark 7.21. \mathcal{F}_R^I , with the lattice structure defined above, constitutes a complete lattice, i.e. every subset of \mathcal{F}_R^I has both a supremum and an infimum.

Now, for every $\psi \in \mathcal{F}_R$, we turn to define the function on \mathcal{F}_R^I , whose greatest fixed point we denote by $C_I^\delta(\psi)$.

Definition 7.22 (δ -Common Knowledge). *Let (γ, R, I, δ) be a δ -coordination-spec. For each $\psi \in \mathcal{F}_R$, we define*

$$f_\psi^\delta : \mathcal{F}_R^I \rightarrow \mathcal{F}_R^I$$

$$(x_i)_{i \in I} \mapsto \left(\bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)} (K_j(\psi \cap x_j)) \right)_{i \in I}, *$$

and denote its greatest fixed point by $C_I^\delta(\psi)$ (“ δ -common knowledge” of ψ by I).

We now show that $C_I^\delta(\psi)$ is well defined. Furthermore, we prove some basic properties of $C_I^\delta(\psi)$, as well as of C_I^δ , which constitutes a function from events $\psi \in \mathcal{F}_R$ to I -tuples of events $\bar{\varphi} \in \mathcal{F}_R^I$.

Lemma 7.23. *Let (γ, R, I, δ) be a δ -coordination-spec and let $\psi \in \mathcal{F}_R$.*

1. $C_I^\delta(\psi)$ is well defined, i.e. f_ψ^δ has a greatest fixed point.

2. Let $\bar{\xi} \in \mathcal{F}_R^I$. If $\bar{\xi} \leq f_\psi^\delta(\bar{\xi})$, then $\bar{\xi} \leq C_I^\delta(\psi)$.

3. C_I^δ is monotone: $\psi \subseteq \phi \Rightarrow C_I^\delta(\psi) \leq C_I^\delta(\phi)$ for every two events $\psi, \phi \in \mathcal{F}_R$.

* This definition may be usefully generalized by allowing ψ to depend on j and even on i as well. Our results concerning this generalization are outside the scope of this work.

Proof. By monotonicity of K_i for every $i \in \mathbb{I}_\gamma$ and by monotonicity of $\otimes^{\leq \varepsilon}$ for every $\varepsilon \in \Delta$, we obtain that f_ψ^δ is monotone. By Remark 7.21, and by Tarski's fixed point theorem[22], the set of fixed points of f_ψ^δ has a greatest element, which equals $\bigvee \{\bar{\xi} \in \mathcal{F}_R^I \mid \bar{\xi} \leq f_\psi^\delta(\bar{\xi})\}$. This proves the first two parts of the lemma.

To prove monotonicity of C_I^δ , let $\psi, \phi \in \mathcal{F}_R$ s.t. $\psi \subseteq \phi$. Once again, by monotonicity of K_i for every $i \in \mathbb{I}_\gamma$ and by monotonicity of $\otimes^{\leq \varepsilon}$ for every $\varepsilon \in \Delta$, we obtain that $f_\psi^\delta(\bar{\varphi}) \leq f_\phi^\delta(\bar{\varphi})$ for every $\bar{\varphi} \in \mathcal{F}_R^I$. By substituting $\bar{\varphi} \triangleq C_I^\delta(\phi)$, and by definition of C_I^δ , we obtain: $C_I^\delta(\psi) = f_\psi^\delta(C_I^\delta(\psi)) \leq f_\phi^\delta(C_I^\delta(\psi))$. By directly applying the second part of the lemma, we obtain that $C_I^\delta(\psi) \leq C_I^\delta(\phi)$. \square

It is now time to prove an equivalent of Theorems 7.10, 7.14 and 7.16, for δ -common knowledge.

Theorem 7.24. *Let (γ, R, I, δ) be a δ -coordination-spec.*

1. *For every event ψ , the ensemble $(K_i(C_I^\delta(\psi)_i))_{i \in I}$ is δ -coordinated.*
2. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is a δ -coordinated ensemble, then $\mathbf{e}_i \subseteq K_i(C_I^\delta(\cup \bar{\mathbf{e}})_i)$ for every $i \in I$.*
3. *If $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ is a δ -coordinated ensemble, then $\cup \bar{\mathbf{e}} \subseteq \cup C_I^\delta(\cup \bar{\mathbf{e}})$.*

Proof. We begin the proof of the first part by noting that by Remark 7.7, $\bar{\mathbf{e}} \triangleq (K_i(C_I^\delta(\psi)_i))_{i \in I}$ is indeed an ensemble for I . Let $(i, j) \in I^2$ and let $(r, t) \in \mathbf{e}_i$. By definition of \mathbf{e}_i and by the truth axiom for knowledge, $(r, t) \in C_I^\delta(\psi)_i$. By definition of C_I^δ ,

$$C_I^\delta(\psi)_i = \bigcap_{k \in I \setminus \{i\}} \otimes^{\leq \delta(i, k)} (K_k(\psi \cap \mathbf{e}_k)) \subseteq \otimes^{\leq \delta(i, j)} (K_j(\psi \cap \mathbf{e}_j)).$$

Thus, we obtain $(r, t) \subseteq \otimes^{\leq \delta(i, j)} (K_j(\psi \cap \mathbf{e}_j))$. By definition of $\otimes^{\leq \delta(i, j)}$, there exists $t' \in \mathbb{T}$ s.t. $t' \leq t + \delta(i, j)$ and $(r, t') \in K_j(\psi \cap \mathbf{e}_j)$. By monotonicity of K_j and by locality of \mathbf{e}_j to j , we obtain $(r, t') \in K_j(\mathbf{e}_j) = \mathbf{e}_j$, and the proof of the first part is complete.*

We move on to proving the second part. Let $\bar{\mathbf{e}}$ be as defined in this part of the theorem. First, we show that $\bar{\mathbf{e}} \leq f_{\cup \bar{\mathbf{e}}}^\delta(\bar{\mathbf{e}})$. Let $i \in I$. Let $(r, t) \in \mathbf{e}_i$ and let $j \in I \setminus \{i\}$. Since $\bar{\mathbf{e}}$ is δ -coordinated, there exists $t' \in \mathbb{T}$ s.t. $t' \leq t + \delta(i, j)$ and $(r, t') \in \mathbf{e}_j$.

* The attentive reader may notice a conceptual similarity between the above argument and the proof of Corollary 6.35. Furthermore, as noted there, this similarly extends to the second part of the proof of Lemma 5.7 as well.

By definition of an ensemble, \mathbf{e}_j is local to j , and thus $\mathbf{e}_j = K_j(\mathbf{e}_j)$. Therefore, $(r, t') \in K_j(\mathbf{e}_j)$. By definition of $\odot^{\leq \delta(i,j)}$, we obtain $(r, t) \in \odot^{\leq \delta(i,j)}(K_j(\mathbf{e}_j))$. Thus,

$$\mathbf{e}_i \subseteq \bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)}(K_j(\mathbf{e}_j)) = \bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)}(K_j((\cup \bar{\mathbf{e}}) \cap \mathbf{e}_j)) = f_{\cup \bar{\mathbf{e}}}^\delta(\bar{\mathbf{e}})_i.$$

By the second part of Lemma 7.23, we thus have $\bar{\mathbf{e}} \leq C_I^\delta(\cup \bar{\mathbf{e}})$. For every $i \in I$, by monotonicity of K_i we obtain $K_i(\mathbf{e}_i) \subseteq K_i(C_I^\delta(\cup \bar{\mathbf{e}})_i)$, and by locality of \mathbf{e}_i to i , we complete the proof of the second part of Theorem 7.24, as $\mathbf{e}_i = K_i(\mathbf{e}_i)$. Let $i \in I$. As we have just shown that $\mathbf{e}_i \subseteq C_I^\delta(\cup \bar{\mathbf{e}})_i$, we also have $\mathbf{e}_i \subseteq \cup C_I^\delta(\cup \bar{\mathbf{e}})$. As this holds for every $i \in I$, and as the r.h.s. does not depend on i , we obtain $\cup \bar{\mathbf{e}} \subseteq \cup C_I^\delta(\cup \bar{\mathbf{e}})$, completing the third, and last, part of the proof. \square

Theorem 7.24, which we have just proved, provides us with some key properties of δ -common knowledge: The first part of the theorem says that the ensemble defined by it is δ -coordinated. The second part says that regardless of the way a δ -coordinated ensemble is formed (be it using δ -common knowledge of some event ψ , or otherwise), the fact that its i 'th coordinate holds implies that i knows the i 'th coordinate of δ -common knowledge of (the disjunction of) this ensemble. The third part, similarly, says that in this case the fact that any coordinate of such an ensemble holds implies that at least one coordinate of δ -common knowledge of (the disjunction of) this ensemble holds.

While Theorem 7.24 does indeed provide us with several key properties of δ -common knowledge, a second, deeper look at this theorem (resp. at its analogues from [11, Section 11.6] surveyed in the previous section) reveals that it does not characterise δ -common knowledge (resp. common knowledge, eventual common knowledge, or ε -common knowledge). Indeed, this theorem (resp. all its analogues) would still hold if we defined each coordinate of δ -common knowledge (resp. common knowledge, eventual common knowledge, or ε -common knowledge) simply as \mathcal{F}_R , i.e. the event “True”.

We remark that it may be verified that the ensemble defined by common knowledge of an event ψ is the greatest perfectly coordinated ensemble $\bar{\mathbf{e}}$ satisfying $\cup \bar{\mathbf{e}} \subseteq \psi$.^{*} Similarly, we remark without a proof (as a proof would be similar to

^{*} Greatest, here, is in the sense of the lattice structure defined in Definition 7.20. In particular, this characterisation implies that such a greatest ensemble exists. This characterisation of the ensemble defined by common knowledge is an immediate consequence of Aumann's definition[3] of common knowledge, which may be rephrased as follows: $C_I(\psi)$, for $I \subseteq \mathbb{I}_\gamma$ and $\psi \in \mathcal{F}_R$, is the greatest event $\phi \subseteq \psi$ that is local to each $i \in I$.

that of Claim 7.27 below) that it may be verified that the ensemble defined by eventual common knowledge of an event of the form $\Diamond(\psi)$ is the greatest eventually-coordinated ensemble \bar{e} satisfying $\cup \bar{e} \subseteq \Diamond(\psi)$. Analogous characterisations, for the ensembles defined by ε -common knowledge and by δ -common knowledge, are, however, more elusive to phrase (as is an analogous characterisation of the ensemble defined by eventual-common knowledge of an arbitrary event.) For this reason, we now only characterise the ensemble defined by δ -common knowledge of events that we call “atemporal”. (It may be readily verified, along the same lines, that an analogous characterisation for the ensemble defined by ε -common knowledge of an atemporal event holds as well.) While this characterisation is similar to that of the ensemble defined by eventual-common knowledge, which we phrased above, it is not analogous in that it conceptually does a significantly less adequate job in capturing the essence of δ -common knowledge (and of ε -common knowledge). Nonetheless, this characterisation suffices for our analysis of the timely-coordinated response problem in the next section.

Definition 7.25 (Atemporal Event). *Let γ be a context and let $R \subseteq \mathcal{R}_\gamma$. We call an event $\psi \in \mathcal{F}_R$ “atemporal” if ψ holding at some time during a run implies that it holds at all times throughout that run. Formally, ψ is atemporal iff it is of the form $R' \times \mathbb{T}$ for some $R' \subseteq R$.*

Remark 7.26. *By Definition 7.25:*

- ψ is atemporal iff $\psi = \Diamond(\psi)$.
- By Remark 7.19 (additivity), $\Diamond = \odot^{\leq \infty}$ is idempotent. Thus, $\Diamond(\psi)$ is atemporal for every $\psi \in \mathcal{F}_R$.

We conclude this section with the following claim, which, together with Remark 7.26, provides a characterisation of the ensemble defined by δ -common knowledge of an atemporal event ψ : It is the greatest δ -coordinated ensemble \bar{e} satisfying $\cup \bar{e} \subseteq \psi$. (In the next section, we conclude that this characterisation holds also for the ensemble defined by δ -common knowledge of a temporal event of the form $\odot^{\leq 0}(\tilde{e})$, where \tilde{e} is an ND event.)

Claim 7.27. *Let (γ, R, I, δ) be a δ -coordination-spec, let $\psi \in \mathcal{F}_R$ and let $\bar{e} \in \mathcal{F}_R^I$ be the ensemble defined by $\mathbf{e}_i \triangleq K_i(C_I^\delta(\psi)_i)$ for every $i \in I$.*

1. $\bar{e}' \leq \bar{e}$, for every δ -coordinated ensemble $\bar{e}' \in \mathcal{F}_R^I$ satisfying $\cup \bar{e}' \subseteq \psi$.

2. $\cup \bar{\mathbf{e}} \subseteq \Diamond(\psi)$. (In particular, for atemporal ψ , by Remark 7.26, $\cup \bar{\mathbf{e}} \subseteq \psi$.)

Proof. We begin by proving the first part. We have

$$\begin{aligned}
\mathbf{e}'_i &\subseteq && \text{by the second part of Theorem 7.24} \\
&\subseteq K_i(C_I^\delta(\cup \bar{\mathbf{e}}')_i) \subseteq && \text{by monotonicity of } K_i \text{ and of } C_I^\delta \\
&\subseteq K_i(C_I^\delta(\psi)_i) = && \text{by definition of } \mathbf{e}_i \\
&= \mathbf{e}_i,
\end{aligned}$$

as required.

To prove the second part, let $i \in I$ and let $j \in I \setminus \{i\}$. By monotonicity of $\odot^{\leq \delta(i,j)}$, we have

$$\begin{aligned}
C_I^\delta(\psi)_i &\subseteq && \text{by definition of } C_I^\delta \\
&\subseteq \odot^{\leq \delta(i,j)} (K_j(\psi \cap C_I^\delta(\psi)_j)) \subseteq && \text{by monotonicity of } K_j \\
&\subseteq \odot^{\leq \delta(i,j)} (K_j(\psi)) \subseteq && \text{by the truth axiom for knowledge} \\
&\subseteq \odot^{\leq \delta(i,j)} (\psi) \subseteq && \text{by Remark 7.19 (monotonicity)} \\
&\subseteq \Diamond(\psi).
\end{aligned}$$

By monotonicity of K_i and by the truth axiom for knowledge, we obtain that $\mathbf{e}_i \subseteq K_i(\Diamond(\psi)) \subseteq \Diamond(\psi)$, completing the proof of the second part. \square

7.3 Analyzing Timely-Coordinated Response

We now relate the machinery developed in the previous section to the timely-coordinated response problem. We begin by formally introducing external inputs as events in \mathcal{F}_R , and by formalizing the relationship between ND events and knowledge.

Definition 7.28. *Given a context γ and an external input $\tilde{e} \in \tilde{E}_\gamma$, we formally associate \tilde{e} with the event “ \tilde{e} is occurring (right now)”, i.e. with the set of all points at which \tilde{e} occurs.*

Remark 7.29. *As noted in Chapter 3, since \tilde{e} is an ND event, it cannot be foreseen (i.e. known to occur) by any agent before it occurs. In the notation of this chapter, this may be formalized as follows: $K_i(\Diamond(\tilde{e})) \subseteq \odot^{\leq 0}(\tilde{e})$ for every $i \in \mathbb{I}_\gamma$. (This follows*

straight from applying Claim 6.27 to the first part of Remark 6.26, at $t_{\tilde{e}}^*$.)

Corollary 7.30. *Let γ be a context and let $\tilde{e} \in \tilde{E}_\gamma$.*

1. $K_i(\Diamond(\tilde{e})) = K_i(\odot^{\leq 0}(\tilde{e}))$, for every $i \in \mathbb{I}_\gamma$.
2. Let $I \subseteq \mathbb{I}_\gamma$ and let $\bar{\mathbf{e}} \in \mathcal{F}_R^I$ be an ensemble. If $\cup \bar{\mathbf{e}} \subseteq \Diamond(\tilde{e})$, then $\cup \bar{\mathbf{e}} \subseteq \odot^{\leq 0}(\tilde{e})$.

Proof. Let $i \in \mathbb{I}_\gamma$. By Remark 7.19 (monotonicity) and by monotonicity of K_i , we have $K_i(\odot^{\leq 0}(\tilde{e})) \subseteq K_i(\Diamond(\tilde{e}))$. Conversely, by the positive introspection axiom, by monotonicity of K_i and by Remark 7.29, we have $K_i(\Diamond(\tilde{e})) = K_i(K_i(\Diamond(\tilde{e}))) \subseteq K_i(\odot^{\leq 0}(\tilde{e}))$, and the proof of the first part is complete.

To prove the second part, let $i \in I$. By locality of \mathbf{e}_i and by monotonicity of K_i , we have $\mathbf{e}_i = K_i(\mathbf{e}_i) \subseteq K_i(\Diamond(\tilde{e}))$. By Remark 7.29, the proof is complete. \square

We conclude this chapter by applying the machinery developed throughout it to obtain an optimal response logic for the timely-coordinated response problem.

Corollary 7.31. *Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec. An optimal response logic for solving $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ is, for every $i \in I$: “respond when $K_i(C_I^\delta(\odot^{\leq 0}(\tilde{e}))_i)$ holds for the first time”.*

Proof. Assume that $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ is solvable and let $P \in TCR_\gamma\langle \tilde{e}, I, \delta \rangle$. W.l.o.g., P is a full-information protocol.[§] For every $i \in I$, define $\mathbf{r}_i \in \mathcal{F}_{R_\gamma(P)}$ as the set of all points at which i responds according to P . Since P is a full-information protocol, the actions (and in particular, the responses) of each agent $i \in I$ at each time $t \in \mathbb{T}$ may be deduced from its state at t . Therefore, $\bar{\mathbf{r}}$ is an ensemble. In addition, define an ensemble $\bar{\mathbf{e}} \in \mathcal{F}_{R_\gamma(P)}^I$ by $\mathbf{e}_i \triangleq K_i(C_I^\delta(\odot^{\leq 0}(\tilde{e}))_i)$. (This is indeed an ensemble, by Remark 7.7, and thus i may indeed respond according to it.)

* For the continuous time model presented in Appendix A, our analyses from both this and the previous chapter depend, through Claim 6.27, on the “no foresight” property. The dependence of the analysis of this chapter on it, though, is not fundamental, in the sense that it may be readily dropped by replacing $\odot^{\leq 0}(\tilde{e})$ with $\Diamond(\tilde{e})$ in Corollary 7.31. (In contrast, it is not clear that the results of the previous chapter may be easily modified to hold in the absence this property.) Moreover, unlike the analysis of the previous chapter, which heavily relies on the “no extrasensory perception” property of the continuous-time model, the analysis of this chapter does not rely on it at all. (This is not surprising given the fact that this analysis is agnostic to the methods of information gain by agents. Indeed, not even once do we mention messages in this analysis.) These are both examples of the ability of the higher-level approach of this chapter to mask the details of the model in question by phrasing its results in terms of knowledge. (The meaning that the knowledge operator takes on in a specific model, though, depends of course on such properties.)

§ A much weaker assumption regarding P suffices as well, actually.

We first prove the optimality of responding according to $\bar{\mathbf{e}}$. Let $i \in I$. We have to show that for every $r \in R_\gamma^\tilde{\mathbf{e}}(P)$, i would respond in r , according to the response logic defined above, no later than $t_r(i)$ (the response time of i in r according to P). Formally, this amounts to showing that $\mathbf{r}_i \subseteq \odot^{\leq 0}(\mathbf{e}_i)$. Since $P \in TCR_\gamma\langle\tilde{\mathbf{e}}, I, \delta\rangle$, we have $\cup \bar{\mathbf{r}} \subseteq \diamond(\tilde{\mathbf{e}})$ and by the second part of Corollary 7.30, $\cup \bar{\mathbf{r}} \subseteq \odot^{\leq 0}(\tilde{\mathbf{e}})$. Thus, by the first part of Claim 7.27, $\bar{\mathbf{r}} \leq \bar{\mathbf{e}}$. (This may seem like a slightly stronger statement than the required $\forall i \in I : \mathbf{r}_i \subseteq \odot^{\leq 0}(\mathbf{e}_i)$, however, we will show in Corollary 8.8 that these are in fact equivalent for full-information protocols.)

We now prove that responding according to $\bar{\mathbf{e}}$ solves $TCR_\gamma\langle\tilde{\mathbf{e}}, I, \delta\rangle$. By the first part of Theorem 7.24, $\bar{\mathbf{e}}$ is δ -coordinated. Let $i \in I$. Correctness of P implies that \mathbf{r}_i holds at some time along any $r \in R_\gamma^\tilde{\mathbf{e}}(P)$. Using the notation of this chapter, this is formulated as $\diamond(\tilde{\mathbf{e}}) \subseteq \diamond(\mathbf{r}_i)$. Therefore, by the first part of this proof and by monotonicity of \diamond , we obtain $\diamond(\tilde{\mathbf{e}}) \subseteq \diamond(\mathbf{e}_i)$ as well. To complete the proof, we note that by the second part of Claim 7.27 and by Remark 7.19 (additivity), $\cup \bar{\mathbf{e}} \subseteq \diamond(\odot^{\leq 0}(\tilde{\mathbf{e}})) = \diamond(\tilde{\mathbf{e}})$, i.e. $\tilde{\mathbf{e}}$ occurs at some time along every run during which any coordinate of $\bar{\mathbf{e}}$ holds.* \square

* By Corollary 7.30, we have $\cup \bar{\mathbf{e}} \subseteq \odot^{\leq 0}(\tilde{\mathbf{e}})$ as well, explaining our previous statement that Corollary 7.30 proves Remark 4.2.

Chapter 8

The Equivalence of Both Approaches

Corollaries 6.35 and 7.31 both present optimal response logics for solving the timely-coordinated response problem presented in Chapter 4. An obvious consequence is that the response logics defined in both corollaries must somehow be equivalent, at least for full-information protocols, which are assumed in our proofs of these corollaries,* and in a shared-clock model, which is assumed in our proof of Corollary 6.35. In the next chapter, we apply our results from the previous chapters to obtain specialized, somewhat more practical, versions of these results for some naturally-occurring models. However, before starting to do so, we prove the equivalence of Corollaries 6.35 and 7.31 in a somewhat more constructive manner in this chapter, which also sheds some more light on the fixed-point analysis of the previous chapter, and makes the notion of δ -common knowledge more concrete. Our aim is to prove the following result, by searching for a more constructive (yet equivalent) definition of δ -common-knowledge, along the lines of the nested-knowledge definition of common knowledge given in Definition 7.5, rather than those of its fixed-point definition given in Theorem 7.11.

Theorem 8.1. *In a shared-clock model, let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec s.t. $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ is solvable. The response logics for $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ defined in Corollaries 6.35 and 7.31 are equivalent when applied to full-information protocols, i.e. these two response logics yield the exact same responses at the same times in each run.*

* The attentive reader may also notice some similarity in the way the second part of the proof of each of these corollaries makes use of its first part.

In order to prove Theorem 8.1, we perform an analysis of δ -common knowledge of events of the form $\odot^{\leq 0}(\tilde{e})$ in full-information protocols. To make our analysis somewhat cleaner and more generic, we first aim to formally capture the properties of such protocols and of events of the form $\odot^{\leq 0}(\tilde{e})$, which are of interest to us.

8.1 Background

In this section, we review two definitions and some basic properties thereof, from [11, Chapter 4]. We rephrase these to match the notation we have introduced so far.

Definition 8.2 (Stability). *Let γ be a context and let $R \subseteq \mathcal{R}_\gamma$. An event $\psi \in \mathcal{F}_R$ is said to be “stable” if once ψ holds at some time during a run, it continues to hold for the duration of that run. Formally, using our notation, ψ is stable iff $\psi = \odot^{\leq 0}(\psi)$.*

Remark 8.3. *By Definition 8.2:*

- *By Remark 7.19 (additivity), $\odot^{\leq 0}$ is idempotent. Thus, $\odot^{\leq 0}(\psi)$ is a stable event for every $\psi \in \mathcal{F}_R$.*
- *$\psi \cap \phi$ is a stable event for every two stable events $\psi, \phi \in \mathcal{F}_R$.*

Indeed, the property of $\odot^{\leq 0}(\tilde{e})$ that we utilize in this chapter is its stability. We now present the second definition based upon [11, Chapter 4], which we utilize in this chapter.

Definition 8.4 (Perfect Recall). *Let γ be a context. A set of runs $R \subseteq \mathcal{R}_\gamma$ is said to exhibit “perfect recall” if for every $r \in R$, for every $i \in \mathbb{I}_\gamma$ and for every $t, t' \in \mathbb{T}$ s.t. $t' \leq t$, the state of i at t in r uniquely determines the state of i at t' in r .*

Remark 8.5. *$R_\gamma(P)$ exhibits perfect recall for every full-information protocol $P \in \mathbb{P}_\gamma$.*

We now distill the property of full-information protocols that is of interest to us, namely that in sets of runs that exhibit perfect recall (and thus, by Remark 8.5, also in full-information protocols), knowledge of a stable event is itself stable. The following is given in [11, Exercise 4.18(b)], and its proof follows directly from the definitions of stability and of knowledge.

Claim 8.6. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ be a set of runs exhibiting perfect recall and let $\psi \in \mathcal{F}_R$. If ψ is stable, then $K_i(\psi)$ is stable as well, for every $i \in \mathbb{I}_\gamma$.*

8.2 A Constructive Proof

Returning to our results and working toward proving Theorem 8.1, we first derive a stability property for δ -common knowledge.

Claim 8.7. *Let (γ, R, I, δ) be a δ -coordination-spec. For every $\psi \in \mathcal{F}_R$, all coordinates of $C_I^\delta(\psi)$ are stable.*

Proof. Let $i \in I$. By Definition 7.18, it is enough to show that $\odot^{\leq 0}(C_I^\delta(\psi)_i) \subseteq C_I^\delta(\psi)_i$. Indeed, we have

$$\begin{aligned}
 \odot^{\leq 0}(C_I^\delta(\psi)_i) &= && \text{by definition of } C_I^\delta \\
 = \odot^{\leq 0} \left(\bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap C_I^\delta(\psi)_j)) \right) &\subseteq && \text{by Remark 7.19} \\
 \subseteq \bigcap_{j \in I \setminus \{i\}} \odot^{\leq 0}(\odot^{\leq \delta(i,j)}(K_j(\psi \cap C_I^\delta(\psi)_j))) &= && \text{by Remark 7.19 (additivity)} \\
 = \bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap C_I^\delta(\psi)_j)) &= && \text{by definition of } C_I^\delta \\
 = C_I^\delta(\psi)_i. &&& \square
 \end{aligned}$$

It should be noted that stability of δ -common knowledge, as guaranteed by Claim 8.7, does not generally guarantee stability of the ensemble defined by it in the first part of Theorem 7.24. Nonetheless, combining Claims 8.6 and 8.7, we obtain stability of this ensemble in the presence of perfect recall.

Corollary 8.8. *Let (γ, R, I, δ) be a δ -coordination-spec s.t. R exhibits perfect recall. For every $\psi \in \mathcal{F}_R$, all coordinates of the ensemble $(K_i(C_I^\delta(\psi)_i))_{i \in I}$ are stable.*

Claims 8.6 and 8.7 and the proof of Corollary 7.31 lead us to consider, for stable ψ and given perfect recall, a slightly different definition for f_ψ^δ than the one given in Definition 7.22. In order to phrase this definition, we first define, given an event $\psi \in \mathcal{F}_R$, notation standing for the event “ ψ holds at exactly ε time units from now”.

Definition 8.9. *Let γ be a context, let $R \subseteq \mathcal{R}_\gamma$ and let $\varepsilon \in \Delta \setminus \{-\infty, \infty\}$. We define*

$$\begin{aligned}
 \odot^\varepsilon : \mathcal{F}_R &\rightarrow \mathcal{F}_R \\
 \psi &\mapsto \{(r, t) \in \Omega_R \mid (r, t + \varepsilon) \in \psi\}^*
 \end{aligned}$$

* Once again, we use the symbol \odot instead of the standard temporal logic notation \bigcirc , in order to emphasize that ε may be nonpositive.

Remark 8.10. By Definition 8.9, for every event $\psi \in \mathcal{F}_R$ we have:

- $\odot^{\varepsilon_1}(\odot^{\leq \varepsilon_2}(\psi)) = \odot^{\leq \varepsilon_1}(\odot^{\varepsilon_2}(\psi)) = \odot^{\leq \varepsilon_1 + \varepsilon_2}(\psi)$, for every $\varepsilon_1, \varepsilon_2 \in \Delta \setminus \{-\infty, \infty\}$.
- $\odot^\varepsilon(\psi) \subseteq \odot^{\leq \varepsilon}(\psi)$, for every $\varepsilon \in \Delta \setminus \{-\infty, \infty\}$.
- \odot^ε commutes with intersection for every $\varepsilon \in \Delta \setminus \{-\infty, \infty\}$: $\odot^\varepsilon(\cap \Psi) = \cap \{\odot^\varepsilon(\psi) \mid \psi \in \Psi\}$ for every set of events $\Psi \subseteq \mathcal{F}_R$.

We now present our slightly modified definition of f_ψ^δ , which differs from the definition of f_ψ^δ given in Definition 7.22 by the use of $\odot^{\delta(i,j)}$ instead of $\odot^{\leq \delta(i,j)}$, and by intersecting with $\diamond(\psi)$ instead of intersecting over eventual knowledge requirements.*

Definition 8.11. Let (γ, R, I, δ) be a δ -coordination-spec s.t. $\delta > -\infty$. For each $\psi \in \mathcal{F}_R$, we define

$$g_\psi^\delta : \mathcal{F}_R^I \rightarrow \mathcal{F}_R^I$$

$$(x_i)_{i \in I} \mapsto \left(\diamond(\psi) \cap \bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\delta(i,j)}(K_j(\psi \cap x_j)) \right)_{i \in I},$$

and denote its greatest fixed point by $C_I^\delta(\psi)$.

Using an argument completely analogous to the proof of Lemma 7.23, it may be shown that $C_I^\delta(\psi)$ is well defined. Furthermore, the same argument shows that $C_I^\delta(\psi)$ also satisfies the obvious analogues of the second and third parts of Lemma 7.23, with regard to g_ψ^δ .

We now present a key observation, which stands at the heart of our proof of Theorem 8.1. While, even in full-information protocols and when ψ is stable, $g_\psi^\delta \neq f_\psi^\delta$ (e.g. when applied to certain unstable events), it so happens that under certain conditions, the greatest fixed points of both of these functions coincide.

Lemma 8.12. Let (γ, R, I, δ) be a δ -coordination-spec s.t. R exhibits perfect recall and s.t. $\delta > -\infty$, and let $\psi \in \mathcal{F}_R$. If ψ is stable, and if $\diamond(\psi) \subseteq \diamond(C_I^\delta(\psi)_i)$ for every $i \in I$, then $C_I^\delta(\psi) = C_I^\delta(\psi)$.

* The intersection with $\diamond(\psi)$ has any effect only if the intersection following it is empty.

*Proof.** \geq : Let $i \in I$. By Claim 8.7, $C_I^\delta(\psi)_j$ is stable for every $j \in I$. Since ψ is stable as well, Remark 8.3 yields that $\psi \cap C_I^\delta(\psi)_j$ is stable for every $j \in I$. We also note that for every $j \in I \setminus \{i\}$, by the truth axiom for knowledge and by Remark 7.19 (monotonicity), we have

$$\odot^{\leq \delta(i,j)} (K_j(\psi \cap C_I^\delta(\psi)_j)) \subseteq \Diamond(\psi). \quad (8.1)$$

Thus, we obtain

$$\begin{aligned} C_I^\delta(\psi)_i &= && \text{by definition of } C_I^\delta \\ &= \bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)} (K_j(\psi \cap C_I^\delta(\psi)_j)) = && \text{by (8.1)} \\ &= \Diamond(\psi) \cap \bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)} (K_j(\psi \cap C_I^\delta(\psi)_j)) \subseteq && \text{intersecting over fewer events} \\ &\subseteq \Diamond(\psi) \cap \bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \odot^{\leq \delta(i,j)} (K_j(\psi \cap C_I^\delta(\psi)_j)) = && \text{by Remark 8.10} \\ &= \Diamond(\psi) \cap \bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \odot^{\delta(i,j)} (\odot^{\leq 0} (K_j(\psi \cap C_I^\delta(\psi)_j))) = && \begin{array}{l} \text{by Claim 8.6 and by} \\ \text{stability of } \psi \cap C_I^\delta(\psi)_j \end{array} \\ &= \Diamond(\psi) \cap \bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \odot^{\delta(i,j)} (K_j(\psi \cap C_I^\delta(\psi)_j)) = && \text{by definition of } g_\psi^\delta \\ &= g_\psi^\delta(C_I^\delta(\psi))_i. \end{aligned}$$

Thus, by the analogue of the second part of Lemma 7.23 for g_ψ^δ , we obtain $C_I^\delta(\psi) \leq \mathcal{C}_I^\delta(\psi)$, as required.

\leq : For every $i \in I$, we have

$$\begin{aligned} \mathcal{C}_I^\delta(\psi)_i &= && \text{by definition of } \mathcal{C}_I^\delta \\ &= \Diamond(\psi) \cap \bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\delta(i,j)} (K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \subseteq && \text{by Remark 8.10} \end{aligned}$$

* It should be noted that we could have saved ourselves some hardship in this proof by replacing $\bigcap_{j \in I \setminus \{i\}} \dots$ with $\Diamond(\psi) \cap (\bigcap_{j \in I, \delta(i,j) < \infty} \dots)$ when defining f_ψ^δ , which would still have allowed us to obtain Corollary 7.31. While this is indeed true, in this case many of our results regarding δ -common knowledge would have required the additional assumption that $\Diamond(\psi) \subseteq \Diamond(C_I^\delta(\psi)_i)$, reducing from their generality and usefulness. The added strength of the approach we have chosen presents itself both in Corollary 9.16, and while discussing eventual common knowledge in Chapter 10.

$$\begin{aligned}
& \subseteq \diamond(\psi) \cap \bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \subseteq & \text{as } \diamond(\psi) \subseteq \diamond(C_I^\delta(\psi)_i) \\
& \subseteq \diamond \left(\bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap C_I^\delta(\psi)_j)) \right) \cap \\
& \quad \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \right) \subseteq & \text{by monotonicity of } \diamond \\
& \subseteq \diamond \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) = \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap C_I^\delta(\psi)_j)) \right) \cap \\
& \quad \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \right) \subseteq & \text{by Remark 7.19} \\
& \subseteq \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) = \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap C_I^\delta(\psi)_j)) \right) \cap \\
& \quad \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \right) \subseteq & \begin{array}{l} \text{by monotonicity of } \odot^{\leq i,j} \\ \text{and of } K_j, \text{ and by the} \\ \text{first part of this proof} \end{array} \\
& \subseteq \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) = \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \right) \cap \\
& \quad \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) \neq \infty}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) \right) = \\
& = \bigcap_{j \in I \setminus \{i\}} \odot^{\leq \delta(i,j)}(K_j(\psi \cap \mathcal{C}_I^\delta(\psi)_j)) = & \text{by definition of } f_\psi^\delta \\
& = f_\psi^\delta(C_I^\delta(\psi))_i.
\end{aligned}$$

Thus, by the second part of Lemma 7.23, we have $\mathcal{C}_I^\delta(\psi) \leq C_I^\delta(\psi)$. \square

One may wonder why we have worked so hard to obtain δ -common knowledge, under the conditions of this chapter, as a fixed point of g_ψ^δ rather than of f_ψ^δ . The answer is simple: g_ψ^δ commutes with the meet operation, while f_ψ^δ does not. (Moreover, as a result, g_ψ^δ is downward-continuous while f_ψ^δ , even in a discrete-time model, is not.) This fact paves our way toward proving Theorem 8.1.

Proof of Theorem 8.1. Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec s.t. $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ is solvable, and let $P \in \mathbb{P}_\gamma$ be a full-information protocol. In this proof, we work in $\Omega_{R_\gamma(P)}$. Note that by Remark 8.3, $\otimes^{\leq 0}(\tilde{e})$ is stable.

By Remark 4.12, solvability of $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ implies $\delta > -\infty$. Furthermore, as shown in the proof of Corollary 7.31, solvability of $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ implies $\diamond(\tilde{e}) \subseteq \diamond(K_i(C_I^\delta(\otimes^{\leq 0}(\tilde{e})))_i)$ for every $i \in I$. By Remark 7.19 (additivity), by the truth axiom for knowledge and by monotonicity of \diamond , we have $\diamond(\otimes^{\leq 0}(\tilde{e})) = \diamond(\tilde{e}) \subseteq \diamond(C_I^\delta(\otimes^{\leq 0}(\tilde{e})))_i$ for every $i \in I$ as well. Thus, by Remark 8.5 and by Lemma 8.12, we obtain $C_I^\delta(\otimes^{\leq 0}(\tilde{e})) = C_I^\delta(\otimes^{\leq 0}(\tilde{e}))$.

It is easy to verify that $g_{\otimes^{\leq 0}(\tilde{e})}^\delta$ commutes with both finite, and infinite, meet. Thus, it is downward-continuous and by Kleene's fixed point theorem*, we obtain

$$C_I^\delta(\otimes^{\leq 0}(\tilde{e})) = \bigwedge_{n \in \mathbb{N}} g_{\otimes^{\leq 0}(\tilde{e})}^{\delta n}(\Omega_{R_\gamma(P)}^I).$$

By \otimes^ε commuting with intersection for every $\varepsilon \in \Delta$, and by K_i commuting with intersection for every $i \in I$, we thus obtain, for every $i \in I$, that

$$\begin{aligned} & C_I^\delta(\otimes^{\leq 0}(\tilde{e}))_i = \\ &= \bigcap_{n \in \mathbb{N}} g_{\otimes^{\leq 0}(\tilde{e})}^{\delta n}(\Omega_{R_\gamma(P)}^I)_i = \\ &= \diamond(\otimes^{\leq 0}(\tilde{e})) \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \otimes^{\delta(i,j)}(K_j(\otimes^{\leq 0}(\tilde{e}))) \right) \cap \\ &\quad \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \otimes^{\delta(i,j)} \left(K_j \left(\otimes^{\leq 0}(\tilde{e}) \cap \bigcap_{\substack{k \in I \setminus \{j\} \\ \delta(j,k) < \infty}} \otimes^{\delta(j,k)}(K_k(\otimes^{\leq 0}(\tilde{e}))) \right) \right) \right) \cap \\ &\quad \cap \dots = \end{aligned}$$

* This fixed point theorem seems to be popularly named after Kleene, as the idea of using the orbit of an extremal element to obtain a fixed point was first used in his proof of his first recursion theorem[15, p. 348]. For a definition of this theorem that is phrased in terms of lattices, continuity and greatest fixed point, see [16].

$$\begin{aligned}
&= \Diamond(\tilde{e}) \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \odot^{\delta(i,j)}(K_j(\odot^{\leq 0}(\tilde{e}))) \right) \cap \\
&\quad \cap \left(\bigcap_{\substack{j \in I \setminus \{i\} \\ \delta(i,j) < \infty}} \odot^{\delta(i,j)} \left(K_j \left(\bigcap_{\substack{k \in I \setminus \{j\} \\ \delta(j,k) < \infty}} \odot^{\delta(j,k)}(K_k(\odot^{\leq 0}(\tilde{e}))) \right) \right) \right) \cap \\
&\quad \cap \dots = \\
&= \Diamond(\tilde{e}) \cap \bigcap_{\substack{\bar{p} \in \mathcal{P}(G_\delta) \\ p_1 = i \\ \bar{p} \neq (i)}} \odot^{\delta(p_1, p_2)}(K_{p_2}(\dots(\odot^{\delta(p_{n-1}, p_n)}(K_{p_n}(\odot^{\leq 0}(\tilde{e})))) \dots)).
\end{aligned}$$

By Corollary 7.30, $K_i(\Diamond(\tilde{e})) = K_i(\odot^{\leq 0}(\tilde{e}))$ and thus, by K_i commuting with intersection, we obtain (omitting henceforth some parentheses for readability)

$$K_i(C_I^\delta(\odot^{\leq 0}(\tilde{e}))_i) = \bigcap_{\substack{\bar{p} \in \mathcal{P}(G_\delta) \\ p_1 = i}} K_{p_1} \odot^{\delta(p_1, p_2)} K_{p_2} \dots \odot^{\delta(p_{n-1}, p_n)} K_{p_n}(\odot^{\leq 0}(\tilde{e})). \quad (8.2)$$

Thus, the response logic from Corollary 7.31 is equivalent, for every $i \in I$, to: “respond as soon as

$$K_{p_1} \odot^{\delta(p_1, p_2)} K_{p_2} \odot^{\delta(p_2, p_3)} \dots K_{p_{n-1}} \odot^{\delta(p_{n-1}, p_n)} K_{p_n}(\odot^{\leq 0}(\tilde{e})) \quad (8.3)$$

holds for every path $\bar{p} \in \mathcal{P}(G_\delta)$ starting at $p_1 = i$.”

By directly applying the methods of Ben-Zvi and Moses[6, 5, 7, 8],* it can be seen that since P is a full-information protocol in a shared-clock model, a (\bar{p}, δ) -traversing \tilde{e} -centipede by $t \in \mathbb{T}$ in a run of P is equivalent to the following holding during that run, expressed by means of their absolute-time modal-logic notation from [8]:

$$K_{(p_1, t_1)} K_{(p_2, t_2)} \dots K_{(p_n, t_n)} \tilde{e}, \quad (8.4)$$

for $t_k \triangleq t + L_{G_\delta}((p_m)_{m=1}^k)$ for every $k \in [n]$, and where \tilde{e} is a proposition corresponding to our $\odot^{\leq 0}(\tilde{e})$ event.

* Ben-Zvi and Moses show this in a discrete-time model. However, it may be verified that the “no foresight” and “no extrasensory perception” properties of the continuous-time model presented in Appendix A suffice in order to adapt their argument, without fundamental change, to this model as well.

As (8.3) holding at t is a different notation for (8.4), the proof is complete. \square

We conclude this chapter with an observation. If $|I| < \infty$ and if G_δ has only trivial (i.e. singleton) strongly connected components, then there are only finitely many paths in G_δ . In this case, Theorem 6.34 and Corollary 6.35 imply that a timely-coordinated response hinges on only finitely many path-traversing centipedes. (This is indeed the case for the ordered response and weakly-timed response problems studied by Ben-Zvi and Moses, as is shown in Chapter 10.) This observation may seem, at first glance, to clash with the infinite nature of fixed points in general, and of greatest fixed points in particular. It is worthwhile to note that what reconciles these is that in this case, $g_\psi^{\delta|I|}$ is constant and therefore its value, which is a finite intersection of nested-knowledge events, is its only fixed point, and thus its greatest fixed point. Furthermore, by Corollary 7.31, solvability of $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ implies that $\Diamond(\tilde{e}) \subseteq \Diamond(C_I^\delta(\odot^{\leq 0}(\tilde{e})))$ and thus, as noted above, we would still have obtained Corollary 7.31 had we defined f_ψ^δ similarly to g_ψ^δ , but using $\odot^{\leq \delta(i,j)}$ instead of $\odot^{\delta(i,j)}$. In this case, the function $f_\psi^{\delta|I|}$ would have also been constant, and the above insight would have held for it as well.

Chapter 9

Results for Practical Models

The analysis of the timely-coordinated response problem in Chapters 6 and 7 is a general one, assuming very little regarding the model in which we work. The advantage of such a general analysis is that the results it yields hold for a vast variety of models and situations. One disadvantage, which we noted in the discussion concluding Chapter 6, is that the gap between these results and their consequences for practical situations is quite large. In this chapter, we derive, from the general analysis of the timely-coordinated response problem from the previous chapters, various stronger results for some special, yet naturally-occurring, cases that we introduce below. We also discuss some possible practical applications of our observations.

9.1 Bounded-Syncausal-Path Contexts

When presenting Theorems 6.12, 6.16 and 6.22 above, we noted that the proofs that Ben-Zvi and Moses present for them strongly rely on time being modeled discretely. It is worthwhile, in this context, to recall the classic “coordinated attack”, or “two generals”, problem[1, 13]. This problem describes a hypothetical situation, in which two army generals, each camped on top a different hill overlooking some village, wish to coordinate a simultaneous attack of this village (i.e. reach common knowledge of an attack time that was not agreed upon in advance), by communicating solely via messengers. While this problem is unsolvable in a discrete-time model[1], Fagin et al.[11, p. 386] note that even in the lack of any delivery guarantee, the generals may successfully coordinate a simultaneous attack if they have access to a messenger who can make infinitely many trips between one general’s camp and the other’s in finite time, by doubling her speed each time she reaches one of the camps.

In fact, Theorem 6.34 implies that in our continuous-time model, even in the absence of any bound guarantee between two disjoint sets of agents I, J (i.e. when $\hat{\delta}_{\mathcal{G}_\gamma}|_{\{I \times J\} \cup \{J \times I\}} \equiv \infty$), a simultaneous response of two agents $i \in I$ and $j \in J$ may be achieved even if no such “infinite syncausal paths” from \tilde{e} to each agent at the time of its response exist, as long as such paths that alternate between these sets of agents an arbitrarily large number of times exist.* In the above scenario, this means that the generals may also coordinate a simultaneous attack if they have access to infinitely many messengers, such that in a finite time frame, for any arbitrarily large $N \in \mathbb{N}$, there exists a messenger who alternates between their camps at least N times. In Corollary 9.16, we formalize the intuition that there are no other ways in which these generals may coordinate even an approximately simultaneous attack.

In order to have any hope of generalizing Theorems 6.12, 6.16 and 6.22 for the timely-coordinated response problem in a continuous-time model, we therefore first have to define some restriction that prevents such Zeno-paradoxical situations from taking place. By doing so, we effectively force the infinitely many associated path-traversing centipedes to degenerate to a broom-like, or centibroom-like, structure. This intuition is formalized in Definition 9.2.

Definition 9.1. *Given a context γ , a run $r \in \mathcal{R}_\gamma$ and two agent-time pairs $\theta_1, \theta_2 \in \mathbb{I}_\gamma \times \mathbb{T}$, we denote by $L_r(\theta_1, \theta_2)$ the supremum of the number of ND events in a syncausal path $\theta_1 \xrightarrow[r]{\gamma} \theta_2$. If $\theta_1 \not\xrightarrow[r]{\gamma} \theta_2$, then we define $L_r(\theta_1, \theta_2) \triangleq \infty$.*

Definition 9.2 (Bounded-Syncausal-Path Context). *We say that a context γ is a “bounded-syncausal-path” context if $\theta_1 \xrightarrow[r]{\gamma} \theta_2$ implies $L_r(\theta_1, \theta_2) < \infty$, for every run $r \in \mathcal{R}_\gamma$ and every two agent-time pairs $\theta_1, \theta_2 \in \mathbb{I}_\gamma \times \mathbb{T}$.*

Remark 9.3. *Some naturally-occurring bounded-syncausal-path contexts include contexts with the following properties, which are customarily taken as axioms:*

- *Any context in which a universal positive lower bound on all delivery times holds, including all contexts of the discrete-time model presented in Chapter 3.*
- *Any context in which only finitely many messages may be sent (or rather, may be delivered early) in any bounded time frame.*

While, due to the continuous nature of time and to the possibility of infinitely many agents, finite-memory and finite-processing-power models need not guarantee

* This is thus also possible in the lack of any delivery guarantee.

bounded-syncausal-path contexts (at least when dealing with protocols that are not necessarily full-information ones), we show in the next section that many results that hold for bounded-syncausal-path contexts, still hold when the memory or processing power of each agent is limited.

We are now ready to generalize Theorems 6.12 and 6.22 and the proofs of Ben-Zvi and Moses[6, 5, 8] for these theorems, for the timely-coordinated response problem in a continuous-time, yet bounded-syncausal-path, context. The following is, in a sense, a converse of Remark 6.36 for such contexts.

Corollary 9.4. *Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec, let $P \in TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ and let $r \in R_\gamma^\tilde{e}(P)$. Let $J \subseteq I$ be a finite subset of I that is contained entirely within one strongly-connected component of G_δ .*

1. *If there exists $i \in J$ s.t. $L_r(\tilde{e}, (i, t_r(i))) < \infty$,* then there exists an event $e \in PND_\gamma^r(i, t_r(i))$ that is an \tilde{e} -broom for J in r . Furthermore, the horizon of this broom may be bounded by a finite bound of the form*

$$\tilde{b}(t_r(i), \max(\hat{\delta}|_{J^2}), |J|, L_r(\tilde{e}, (i, t_r(i)))).$$

2. *If, in addition, $\hat{\delta}|_{J^2}$ is antisymmetric, then the broom guaranteed by the first part of this corollary is by $(t_r(j))_{j \in J}$, implying the results of Theorems 6.12 and 6.22 for bounded-syncausal-path contexts in a continuous-time model.*

Proof. Denote $n \triangleq |J|$. The fact that J resides within one strongly-connected component of G_δ implies $\hat{\delta}|_{J^2} < \infty$, and thus, a Hamiltonian cycle exists in the subgraph of G_δ induced by J . Let (p_1, \dots, p_n, p_1) be such a cycle, for which $p_1 = i$. We now concatenate this cycle to itself enough times to obtain a path of $l \triangleq (n-1) \cdot L_r(\tilde{e}, (i, t_r(i))) + 1$ vertices, which we denote by

$$\bar{p}' = \underbrace{(p_1, \dots, p_n, p_1, \dots, p_n, p_1, \dots, p_n, p_1, \dots, p_k)}_{l=(n-1) \cdot L_r(\tilde{e}, (i, t_r(i))) + 1}.$$

By Theorem 6.34, r contains a $(\bar{p}', \hat{\delta})$ -traversing \tilde{e} -centipede by $t_r(i)$ — denote it by $\bar{e} = (e_m)_{m=1}^l$. By definition, \bar{e} contains at most $L_r(\tilde{e}, (i, t_r(i)))$ distinct events. Thus, each distinct event contained in \bar{e} appears in it, on average, at least $l/L_r(\tilde{e}, (i, t_r(i))) > n-1$ times. Thus, by the pigeonhole principle, there exists an

* In particular, this holds for every $i \in J$, if γ is a bounded-syncausal-path context.

event $e \in PND_\gamma^r(i, t_r(i))$ that appears in \bar{e} at least n times. By definition of a path-traversing centipede and by antisymmetry of the syncausality relation, these appearances are consecutive, and thus we obtain that there are n consecutive vertices in \bar{p}' to which there exists a delivery guarantee from e . As any set of n consecutive vertices in \bar{p}' exactly equals J , we obtain that e is an \tilde{e} -broom for J in r by $t_r(i) + L_{G_\delta}(\bar{p}')$.

We complete the proof of the first part of the corollary by bounding this time:

$$\begin{aligned}
t_r(i) + L_{G_\delta}(\bar{p}') &\leq && \text{by definition of } L_{G_\delta} \\
&\leq t_r(i) + (l - 1) \cdot \max(\hat{\delta}|_{J^2}) = && \text{by definition of } l \\
&= t_r(i) + (|J| - 1) \cdot L_r(\tilde{e}, (i, t_r(i))) \cdot \max(\hat{\delta}|_{J^2}) < && \text{by finiteness of all elements} \\
&< \infty.
\end{aligned}$$

We note that for large $L_r(\tilde{e}, (i, t_r(i)))$, obtaining the shortest possible \bar{p}' frequently involves choosing a Hamiltonian cycle of minimal length.

We now move on to proving the second part of the corollary. If $\hat{\delta}|_{J^2}$ is antisymmetric, then the length of any $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_{\hat{\delta}|_{J^2}})$ is precisely $\hat{\delta}(p_n) - \hat{\delta}(p_1)$. Therefore, for any $r \in R_\gamma^{\tilde{e}}(P)$, any end node of any $(\bar{p}, \hat{\delta})$ -traversing \tilde{e} -centipede by $t_r(p_1)$ is of the form $(p_k, t_r(p_1) + L_{G_\delta}((p_m)_{m=1}^k)) = (p_k, t_r(p_1) + \hat{\delta}(p_1, p_k)) = (p_k, t_r(p_k))$, and the proof of the second part of the corollary is complete. \square

An analogous proof gives rise to the following corollary, which generalizes both Theorem 6.16 and Corollary 9.4.

Corollary 9.5. *Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec, let $P \in TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ and let $r \in R_\gamma^{\tilde{e}}(P)$. Let $n \in \mathbb{N}$ and let $\bar{J} = (J_m)_{m=1}^n \in (2^I)^n$ be a tuple of finite subsets of I , each of which is contained entirely within one strongly-connected component of G_δ . Assume, furthermore, that no two of these subsets are contained within the same strongly-connected component of G_δ , and that for every $m \in [n - 1]$, there exists a path from J_m to J_{m+1} in G_δ .*

1. *If there exists $i \in J_1$ s.t. $L_r(\tilde{e}, (i, t_r(i))) < \infty$,* then there exists an \tilde{e} -centibroom for \bar{J}^{rev} in r , consisting entirely of events from $PND_\gamma^r(i, t_r(i))$. Furthermore, the horizon of this centibroom may be bounded by a finite bound of the form*

$$\tilde{b}(t_r(i), n, (\max(\hat{\delta}|_{J_m^2}))_{m=1}^n, (\min(\hat{\delta}|_{J_m \times J_{m+1}}))_{m=1}^{n-1}, (|J_m|)_{m=1}^n, L_r(\tilde{e}, (i, t_r(i)))).$$

* As before, this holds for every $i \in J_1$ if γ is a bounded-syncausal-path context.

2. For every $m \in [n]$, choose an arbitrary $j_m \in J_m$. If, in addition to the conditions of the previous part, for every $m \in [n]$, $\hat{\delta}|_{J_m^2}$ is antisymmetric, then the end nodes of the centibroom guaranteed by the first part of this corollary are $\{(j, t_r(j)) \mid j \in J_1\}$, $\{(j, t_r(j_1) + \hat{\delta}(j_1, j)) \mid j \in J_2\}$, $\{(j, t_r(j_1) + \hat{\delta}(j_1, j_2) + \hat{\delta}(j_2, j)) \mid j \in J_3\}$, etc.* In particular, this also implies the result of Theorem 6.16 for bounded-syncausal-path contexts in a continuous-time model.[§]

Proof sketch. For each $m \in [n]$, choose a Hamiltonian cycle in the subgraph of $G_{\hat{\delta}}$ induced by J_m , and concatenate it to itself enough times to obtain a path of $l \triangleq (|J_m| - 1) \cdot L_r(\tilde{e}, i, t_r(i)) + 1$ vertices. Now, apply Theorem 6.34 to the concatenation of all these paths, in ascending order of m . The rest of the proof is analogous to the proof of Corollary 9.4. \square

The proof of Corollary 9.4 gives rise to the following observation, stated using the notation of that proof, and under the assumptions thereof: Every path $\bar{p} \in \mathcal{P}(G_{\hat{\delta}|_{J_2}})$ is a prefix of some path \bar{p}' , for which there exists a (\bar{p}', δ) -traversing \tilde{e} -centipede that has a suffix that constitutes an \tilde{e} -broom for J . Thus, this path-traversing centipede yields a path-traversing centipede for every path of which \bar{p}' is a prefix. While, as noted in the closing remarks of Chapter 6, implementing the optimal response logic from Corollary 6.35 may entail, in the most general setting, checking for infinitely many path-traversing centipedes (using infinitely many facts stored in memory), this observation, and the more general analogous observation stemming from the proof of Corollary 9.5, provide a practical and straightforward approach for implementing this optimal response logic in bounded-syncausal-path contexts, as illustrated in Figure 5, and by the following example.

Example 9.6. Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec s.t. γ is a bounded-syncausal-path context, s.t. $|I| < \infty$ and s.t. $\text{TCR}_{\gamma}(\tilde{e}, I, \delta)$ is solvable. For simplicity of this example, assume for the time being that G_{δ} is strongly connected and that $\delta > 0$. Let $i \in I$

* By antisymmetry of $\hat{\delta}$ on each J_m , these are invariant to the choice of representatives $(j_m)_{m=1}^n$.

[§] The scenario studied in the second part of Corollary 9.5 is, in a sense, a timed generalization of ordered joint response, in that it is to tightly-timed response and to weakly-timed response, as ordered joint response is to simultaneous response and to ordered response. This scenario generalizes all the response problems studied by Ben-Zvi and Moses that are surveyed in Chapter 6. As noted in that chapter, the most general form of this scenario does not fall within the scope of any of the coordinated response problems defined and studied by Ben-Zvi and Moses[6, 5, 7, 8], even though no weak mutual dependencies between response time exist in it.

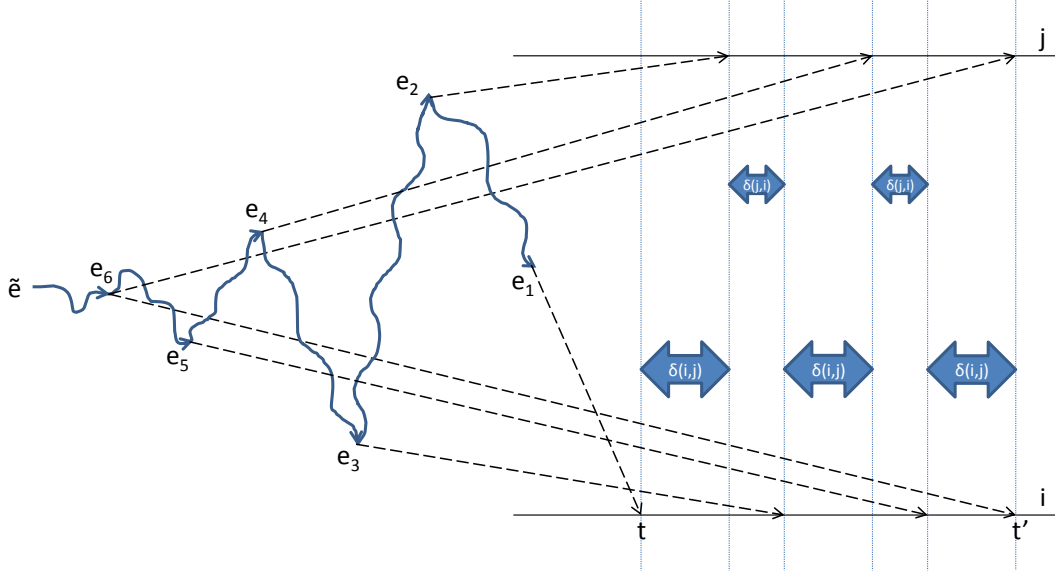


Figure 5: When $I = \{i, j\}$, an $((i, j, i, j, i, j), \delta)$ -traversing \tilde{e} -centipede by t $((e_m)_{m=1}^6)$ with an \tilde{e} -broom suffix by t' (e_6) provides sufficient data for i to respond at t according to the optimal response logic for $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$, as it implies a (\bar{p}, δ) -traversing \tilde{e} -centipede by t for any $\bar{p} \in \mathcal{P}(G_\delta)$ starting at $p_1 = i$.

It should be noted that in this example, common knowledge of the occurrence of \tilde{e} is only attained at t' . As the time at which common knowledge of this occurrence is attainable in a full-information protocol is independent of δ , this intuitively illustrates a property of the timely-coordinated response problem hinted to by the first part of Remark 6.33 and by Example 1.1 opening this work in Chapter 1: intuitively, the greater δ is, the better a chance there is to respond earlier in many cases.

and $t \in \mathbb{T}$. In a full-information protocol, the following algorithm may be applied by i at t to decide whether it should respond at that time according to the optimal response logic for $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$.

The Algorithm: First, check if any $e \in PND_\gamma^r(i, t)$ is an \tilde{e} -broom for I by any past, present, or future time. (This only depends on the observer of e , so we may compute this efficiently with the aid of a precalculated lookup table. Moreover, this may be computed once for each e , storing the result in the state of i .) Denote the set of all such broom events by B .^{*} If $B = \emptyset$, then i should not respond at t .

For each $e \in B$, denote by br_e the earliest time by which e is an \tilde{e} -broom for I . Thus, by locality of bound guarantees and as $\delta > 0$, we obtain that $(e)^n$ is a (\bar{p}, δ) -traversing \tilde{e} -centipede by br_e , and by any later time, for every $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$.

^{*} While $PND_\gamma^r(i, t)$ may be infinite, implying that calculation of B may require infinite processing power and memory, we show in Claim 9.14 in the next section that if $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable by any finite-memory or finite-processing-power protocol, then it is enough to consider only finitely many events from $PND_\gamma^r(i, t)$ at this stage.

(Note that $br_e - t_e$ is a constant depending solely on the observer of e , so once again, some advance computation allows an efficient calculation of br_e , which, once calculated, may be stored in the state of i .)

Denote $\tilde{br} \triangleq \sup_{e \in B} \{br_e\}$. This is a finite quantity, due to $br_e - t_e$ depending only on the observer of e , and by finiteness of I . Denote the set of paths $\bar{p} \in \mathcal{P}(G_\delta)$ s.t. $L_{G_\delta}(\bar{p}) < \tilde{br} - t$ by \mathcal{P} . As $\delta > 0$, and as $|I| < \infty$, we have $|\mathcal{P}| < \infty$ as well.

i should respond at t iff for each $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}$, there exists, by t , a (\bar{p}, δ) -traversing \tilde{e} -centipede $\bar{e} = (e_m)_{m=1}^n \in PND_\gamma^r(i, t)^n$ s.t. there exists $e \in B$ satisfying $e \xrightarrow{\gamma}_r e_n$. (This check may be implemented efficiently using backtracking, and accelerated using some precalculations.)

Dropping Unneeded Assumptions: As noted above, the assumptions that $\delta > 0$ and that G_δ is strongly connected are not required. Handling a situation in which they do not hold is not inherently different, albeit significantly more cumbersome. We now overview the key points of difference regarding these cases.

First, let us drop the assumption that $\delta > 0$. This introduces two obstacles for the above algorithms, which we now rectify.

The first obstacle is that, due to the possibility of δ taking negative values, it may no longer hold that $(e)^n$ is a (\bar{p}, δ) -traversing \tilde{e} -centipede by br_e for every $e \in B$ and $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$. For this to hold again, we redefine br_e , for every $e \in B$, as the earliest time satisfying $e \xrightarrow{\gamma} (i, br_e + \min(\hat{\delta}|_{\{i\} \times I}))$ for every $i \in I$. (See also Remark 6.36.) Note that br_e is finite, by Lemma 5.7. Also note that $br_e - t_e$ still depends only on the observer of e . We accordingly redefine \mathcal{P} as the set of paths $\bar{p} \in \mathcal{P}(G_\delta)$, s.t. the length of \bar{p} , and of every prefix thereof, is less than $\tilde{br} - t$.

The second obstacle is that G_δ may contain nontrivial cycles of zero length, which implies that \mathcal{P} may be of infinite cardinality. The adjustment of the algorithm for this case is somewhat less straightforward. We partition I into pairwise-disjoint equivalence classes s.t. $i, j \in I$ are in the same equivalence class iff $\hat{\delta}(i, j) = -\hat{\delta}(j, i)$. It may be readily verified that these equivalence classes are exactly the subsets of I on which $\hat{\delta}$ is antisymmetric, and that are maximal with regard to this property. Another characterisation of these classes, which is of key importance to us, is that a cycle in G_δ is of length 0 iff all its vertices belong to the same equivalence class. Let $J \subseteq I$ be a set of representatives for all such equivalence classes. Denote, for each $j \in J$, its equivalence class (which it represents) by I_j . We restrict \mathcal{P} to paths containing only “representative” vertices $j \in J$. (Thus, by the above characterisation of equivalence classes using cycles lengths, $|\mathcal{P}| < \infty$ once again.) Finally, for

each $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}$, we require that the matching (\bar{p}, δ) -traversing \tilde{e} -centipede be, in a sense, a “path-traversing centibroom” for I_{p_1}, \dots, I_{p_n} . To be more precise, we require a path-traversing centipede that, for each $k \in [n]$, does not merely satisfy $e_k \xrightarrow{\gamma} (p_k, t + L_{G_\delta}((p_m)_{m=1}^k))$ (as in the path-traversing centipede definition), but also $e_k \xrightarrow{\gamma} (j, t + L_{G_\delta}((p_m)_{m=1}^k) + \hat{\delta}(p_k, j))$, for every $j \in I_{p_k}$. (Checking for the existence of all required path-traversing centipedes/centibrooms may still be efficiently implemented using the same techniques as in the simpler case above.)

Finally, we sketch the key point of adapting the above algorithm for the case in which G_δ is not necessarily strongly connected. In this case, the syncausal structure underlying this algorithm is somewhat more complex. Instead of the algorithm revolving around (\bar{p}, δ) -traversing \tilde{e} -centipedes/centibrooms $(e_m)_{m=1}^n$ for which e_n is an \tilde{e} -broom for all agents in I , the algorithm searches for (\bar{p}, δ) -traversing \tilde{e} -centipedes $(e_m)_{m=1}^n$ for which, for each strongly-connected component I' of G_δ that is visited by \bar{p} , the event $e_{\max\{m \mid p_m \in I'\}}$ may be an \tilde{e} -broom for I' .

When presenting coordinated response problems in Chapter 4, we noted that for many such problems, it is possible to obtain a characterisation for solvability from an optimal response logic. Indeed, Corollary 6.35 implies that a necessary and sufficient condition for solvability of the timely-coordinated response problem in a shared-clock model is a guarantee that in every triggered run, infinitely many path-traversing centipedes (one for each path in G_δ) occur by some finite time. While this indeed fully characterises solvability of the timely-coordinated response problem, the complexity of this characterisation renders it not very usable. Corollary 9.5 also allows us to present a surprisingly simpler characterisation for solvability of the timely-coordinated response problem of finitely many agents in a shared-clock model. We first present a special case thereof, which stems from Corollary 9.4.

Corollary 9.7. *In a shared-clock model, let γ be a context, let $I \subseteq \mathbb{I}_\gamma$ be finite, and let $\tilde{e} \in \tilde{E}_\gamma$. the following conditions are equivalent:*

1. $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ is solvable for some $\delta : I^2 \rightarrow \Delta \setminus \{-\infty\}$ s.t. G_δ is strongly connected and contains no negative cycles.
2. $TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ is solvable for every $\delta : I^2 \rightarrow \Delta \setminus \{-\infty\}$ s.t. G_δ is strongly connected and contains no negative cycles.
3. $SR_\gamma\langle \tilde{e}, I \rangle$ is solvable.

Proof. $3 \Rightarrow 2$: Let $P \in SR_\gamma\langle\tilde{e}, I\rangle$. By Corollary 5.8, δ is implementable. Therefore, by Claim 5.1, $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable.

$2 \Rightarrow 1$: Immediate.

$1 \Rightarrow 3$: Let $P \in TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ for some such δ . W.l.o.g., assume that P is a full-information protocol. By combining the “stand-alone external inputs” and “no foresight” properties of the continuous-time model presented in Appendix A, we may construct a run $r \in R_\gamma^{\tilde{e}}(P)$ in which no message is delivered early less than 1 time unit after it is sent. (For a discrete-time model, this holds for any run $r \in R_\gamma^{\tilde{e}}(P)$.) By Corollary 9.4, r contains an \tilde{e} -broom for I . Thus, P may be modified to solve $SR_\gamma\langle\tilde{e}, I\rangle$ by modifying its response logic for all agents to: “respond at the earliest horizon of an \tilde{e} -broom for all I ”. (The fact that P is a full-information protocol, together with the fact that the clock is shared, guarantees that if a broom for all I exists by any $t \in \mathbb{T}$, then each agent in I can deduce this at t .) \square

Corollary 9.7, together with Theorem 6.12, imply that $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable under condition 1 of this corollary iff there exists an agent $i \in \mathbb{I}_\gamma$ to which there exists a path $\bar{p} \in \mathcal{P}(\mathcal{G}_\gamma)$ from $i_{\tilde{e}}$ and s.t. $\max\{\hat{\delta}_{\mathcal{G}_\gamma}(i, j)\}_{j \in I} < \infty$. Furthermore, if $\delta \geq 0$, then given a full-information protocol endowed with the optimal response logic, the latest of the responses of I occurs, in each triggered run of this protocol, no later than

$$\min\{L_G(\bar{p}) + \max\{\hat{\delta}_{\mathcal{G}_\gamma}(p_n, j)\}_{j \in I} \mid \bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(\mathcal{G}_\gamma) \ \& \ p_1 = i_{\tilde{e}}\}^* \quad (9.1)$$

time units after the occurrence of \tilde{e} , and this bound is tight. (If this value is infinite, then there exist triggered runs in which the time of the latest of the responses is arbitrarily large.)

Using Corollary 9.5, we may similarly deduce a generalization of Corollary 9.7, for the case in which G_δ is not necessarily strongly connected, yielding a characterisation for solvability of any instance of the timely-coordinated response problem of finitely many agents in a shared-clock model.

Corollary 9.8. *In a shared-clock model, let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec s.t. $|I| < \infty$ and s.t. G_δ contains no negative cycles. $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable iff for every tuple (I_1, \dots, I_n) of strongly-connected components of G_δ s.t. there exists a path in G_δ from I_m to I_{m+1} for every $m \in [n-1]$, there exists $(i_m)_{m=1}^n \in \mathbb{I}_\gamma$ satisfying:*

* A generalization of this expression, for cases in which $\delta \not\geq 0$, may be readily obtained by applying Remark 6.36.

- There exists a path in \mathcal{G}_γ from $i_{\tilde{e}}$, through i_n , through i_{n-1}, \dots , to i_1 .
- $\max\{\hat{\delta}_{\mathcal{G}_\gamma}(i_m, j)\}_{j \in I_m} < \infty$, for every $m \in [n]$.

(A natural, yet more cumbersome, analogue of (9.1) may be phrased under the conditions of this corollary as well.)

Corollary 9.8 implies, in particular, that solvability of $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ depends only on the strongly-connected components of G_δ , on the partial order it induces on them, and on \mathcal{G}_γ . In most practical situations, \mathcal{G}_γ is strongly connected, as absence of this property means that there exist two agents $i, j \in \mathbb{I}_\gamma$ s.t. i may never hope to send any data to j , either directly or indirectly. If indeed \mathcal{G}_γ is strongly connected, then Corollary 9.8 reduces to the following, surprisingly simple, condition for solvability of the timely-coordinated response problem, with which we conclude this section.

Corollary 9.9. *In a shared-clock model, let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec s.t. \mathcal{G}_γ is strongly connected, s.t. $|I| < \infty$ and s.t. G_δ contains no negative cycles. The following conditions are equivalent:*

1. $TCR_\gamma\langle\tilde{e}, I, \delta\rangle$ is solvable.
2. $SR_\gamma\langle\tilde{e}, J\rangle$ is solvable, for each strongly-connected component J of G_δ .
3. For every strongly-connected component J of G_δ , there exists $i \in \mathbb{I}_\gamma$ s.t. $\max\{\hat{\delta}_{\mathcal{G}_\gamma}(i, j)\}_{j \in J} < \infty$.

9.2 Finite-Influence Protocols

In this section, we explain, as promised, why the results of the previous section hold also for models in which the memory or processing power of each agent is limited. As a nice bonus, many of those results, which in the previous section held only for finite sets of agent, will turn out to hold in such models for infinite sets of agents as well. To give some intuition for the definition that we use to make this explanation precise, we first prove a result regarding simultaneous response among infinitely many agents.

Recall that Theorem 6.12 implies that in a discrete-time model, simultaneous response of finitely many agents based on an ND event requires the existence of a broom. In Corollary 9.4, we have relaxed the requirement for discrete-time in this

result to a requirement for a bounded-syncausal-path context. A natural question to ask is whether the requirement for finiteness of the set of responding agents may somehow be relaxed as well. A quick check shows that none of the proof strategies we have seen so far for (any variant of) Theorem 6.12 scale to the case of coordinating a simultaneous response among infinitely many agents. Indeed, it turns out that while a broom for any finitely-sized subset of agents is guaranteed under such conditions (it is not hard to show that such a collection of finite brooms is sufficient to optimally coordinate a simultaneous response of all agents),* a broom for all agents is not guaranteed even in bounded-syncausal-path contexts. (As may be expected, such a broom is sufficient for coordinating a simultaneous response, although it is possible to construct an example in which the response logic based on the existence of such a broom is non-optimal.) The following theorem characterises the conditions required for simultaneous response of countable many agents in a bounded-syncausal-path context and shows that there are, in a sense, some unintuitive consequences to the absence of a broom for all agents in this situation. We will shortly use the insights this theorem provides us in order to phrase a restriction on protocols, which disallows such “consequences”.

Theorem 9.10 (Infinite Broom or Infinitely Many Brooms). *Let γ be a bounded-syncausal-path context, let $I \subseteq \mathbb{I}_\gamma$ s.t. $|I| = \aleph_0$, let $\tilde{e} \in \tilde{E}_\gamma$ and let $P \in SR_\gamma\langle\tilde{e}, I\rangle$. For every $r \in R_\gamma^\tilde{e}(P)$, one of the following holds:*

1. *There exists an \tilde{e} -broom for I by t_r in r ; or*
2. *For every finite $J \subseteq I$, there exist infinitely many distinct \tilde{e} -brooms for J by t_r in r .*

Proof. Set $t \triangleq t_r$. Let $J \subseteq I$ be finite. We inductively construct a sequence of sets of ND events $(E_k)_{k=1}^\infty \in ND_\gamma(r)^\mathbb{N}$, satisfying:

* This observation is closely tied to the fact that by Definition 7.5, in any context we have:

$$C_I(\psi) = \bigcap_{\substack{J \subseteq I \\ |J| < \infty}} C_J(\psi),$$

for any set of agents I , regardless of its cardinality. Readers who find it intuitively difficult to accept this “finiteness” of common knowledge (which yields the above mentioned possibility of a simultaneous response of I in a bounded-syncausal-path context in the absence of a broom for I) may wish to read [4], in which Barwise suggests an interpretation of knowledge in which the fixed-point definition of common knowledge (Theorem 7.11) is strictly stronger than its definition as a conjunction of finitely-nested knowledge events (Definition 7.5).

- $E_1 = J$.
- $\forall k \in \mathbb{N} : |E_k| < \infty$.
- $\forall k \in \mathbb{N} : E_k \subseteq E_{k+1}$.
- $\bigcup_{k=1}^{\infty} E_k = I$.

To construct $(E_k)_{k=1}^{\infty}$, choose any well-ordering of I under which it is isomorphic to ω , and given E_k for $k \in \mathbb{N}$, set $E_{k+1} \triangleq E_k \cup \{\min\{I \setminus E_k\}\}$.

For every $k \in \mathbb{N}$, set $B_k \triangleq \{e \in ND_{\gamma}(r) \mid e \text{ is an } \tilde{e}\text{-broom for } E_k \text{ by } t \text{ in } r\}$. By Remark 4.14 and by Corollary 9.4, $\forall k \in \mathbb{N} : B_k \neq \emptyset$. Furthermore, by the definition of a broom and since $(E_k)_{k=1}^{\infty}$ is increasing, $\forall k \in \mathbb{N} : B_k \supseteq B_{k+1}$.

If $|B_1| < \infty$, then $\forall k \in \mathbb{N} : |B_k| < \infty$ and all of $\{B_k\}_{k=1}^{\infty}$ are therefore closed under the co-finite topology on \mathbb{N} and have the finite intersection property. From compactness of \mathbb{N} under this topology, we obtain $B \triangleq \bigcap_{k=1}^{\infty} B_k \neq \emptyset$. Let $e \in B$. By definition of $(B_k)_{k=1}^{\infty}$, e is an \tilde{e} -broom for $\bigcup_{k=1}^{\infty} E_k = I$ by t , and the first condition of Theorem 9.10 is satisfied.

Otherwise, B_1 is of infinite cardinality, and thus J satisfies the second condition of Theorem 9.10. \square

Theorem 9.10 shows that in the absence of a single broom for all agents in I , there are, for each $i \in I$, infinitely many brooms that are “important” to i , in the sense that the existence of any finite subset thereof is not sufficient to trigger the response of i at t . It seems unrealistic for i to check (directly, or indirectly by receiving this information from some other agent) that infinitely many brooms exist.* We now formalize this intuition.

Definition 9.11 (Finite-Influence Protocol). *Let γ be a context. We say that a protocol $P \in \mathbb{P}_{\gamma}$ is a “finite-influence” protocol if for any run $r \in R_{\gamma}(P)$ and for any agent-time pair $(i, t) \in \mathbb{I}_{\gamma} \times \mathbb{T}$ s.t. i responds in r at t , there exists a finite t -retainable set $E \in RND_{\gamma}^P(r, t)$ s.t. i still responds at t in $r \overset{t}{\cap} E$.*

Obviously, all protocols in any context in which only finitely many messages may be sent (or rather, may be delivered early) in any bounded time frame, are

* It should be noted that while one may argue that sending infinitely many messages is equally unrealistic (an argument that suggests that even $ER_{\gamma}(\tilde{e}, I)$ is unsolvable for infinite I), we would like to argue that broadcasting a single message from one agent to infinitely many agents may not be inconceivable. For example, one may post such a message in a publicly-visible place such that each agent is guaranteed to notice this message within a given time period after its posting.

finite-influence protocols. When only finitely many agents exist, then all protocols in any discrete-time context in which a universal positive lower bound on all delivery times holds, are also finite-influence protocols. As noted in the previous section, both of these properties are traditionally taken as axioms, and this is equally true for finiteness of the set of agents. While, as noted above, not all finite-memory and finite-processing power models guarantee bounded-syncausal-path contexts, we now explain why all protocols in such models are finite-influence ones.

If a protocol P is not a finite-influence protocol, then there exists a run $r \in R_\gamma(P)$ and an agent-time pair (i, t) s.t. by time t either i , or some other agent who sends information to i , has to either take infinitely many ND events into account when performing some state change (i.e. no finite subset of these events taking place would have yielded the same state change), or perform infinitely many state changes. This implies that at least one agent utilizes infinite processing power in finite time. If, furthermore, we assume that no two ND events reach an agent at exactly the same time (it is enough to assume that the logic of i can not atomically access information regarding more than one ND event observed by it), then that agent not only utilizes infinite processing power, but also infinite memory, as the infinite amount of state changes described above must involve infinitely many unique states. (Alternatively, the computation of such a single state change that takes into account infinitely many ND events, requires infinite memory.)

We may conclude that for finite-influence protocols in bounded-syncausal-path contexts, Theorem 6.12 holds even when time is continuous and when I may be countably infinite. It took us quite a chain of reductions and conclusions to show this. One may argue that this might suggest that the concept of finite influence is somewhat artificial, and “not from the book”. To try and refute this argument, we now present a novel, direct and concise proof of Theorem 6.12 which holds for any finite-influence protocol (even in a continuous-time model), regardless of the context and of the cardinality of I .*

Theorem 9.12 (Broom). *Let γ be a context, let $I \subseteq \mathbb{I}_\gamma$, let $\tilde{e} \in \tilde{E}_\gamma$ and let $P \in SR_\gamma\langle\tilde{e}, I\rangle$ be a finite-influence protocol. Each $r \in R_\gamma^\tilde{e}(P)$ contains an \tilde{e} -broom for I*

* This proof extends immediately to proving the second part of Corollary 9.4 under the conditions of Claim 9.14 below (implying an infinite-agent version of Theorem 6.22 for finite-influence protocols), and even extends (using an inductive argument very similar to the one used in Theorem 6.34) to prove the second part of Corollary 9.5 under the conditions of Claim 9.14 (implying an infinite-agent version of Theorem 6.16 for finite-influence protocols). For the sake of conciseness, though, we phrase and prove it here only for simultaneous response.

by $t \triangleq t_r$.

Claim 9.13. *Under the conditions of Theorem 9.12, there exists $r' \subseteq^t r$ such that $t_{r'} = t$, such that $|ND_\gamma(r', t)| < \infty$ and such that $PND_\gamma^{r'}(i, t) = ND_\gamma(r', t)$ for every $i \in I$.*

Proof. Let $E \in RND_\gamma^P(r, t)$ be minimal such that $t_{r \cap E} = t$. (There always exists such a finite set, due to P being a finite-influence protocol, since by correctness of P , if even one agent $i \in I$ responds at t , then all agents in I do.) By definition, $E = ND_\gamma(r \overset{t}{\cap} E, t)$.

Let $i \in I$. Set $E' \triangleq PND_\gamma^{r \overset{t}{\cap} E}(i, t)$. By Corollary 6.29, $E' \in RND_\gamma^P(r \overset{t}{\cap} E, t)$, and the state of i at t is the same in $r \overset{t}{\cap} E$ and in $(r \overset{t}{\cap} E) \overset{t}{\cap} E' = r \overset{t}{\cap} E'$. Therefore, i still responds at t in $r \overset{t}{\cap} E'$. Thus, by correctness of P , we obtain $t_{r \overset{t}{\cap} E'} = t$. By minimality of E , therefore, since $E' \subseteq E$, we obtain $E' = E$, and thus any $r' \in r \overset{t}{\cap} E$ fulfills the above requirements. \square

Proof of Theorem 9.12. Let r' be as in Claim 9.13 and choose $e \in ND_\gamma(r', t)$ with maximal t_e among all those satisfying $\tilde{e} \xrightarrow{\gamma_r} e$. (There always exists such an event, by finiteness of $ND_\gamma(r', t)$ and as \tilde{e} is always a viable candidate, since $\tilde{e} \in ND_\gamma(r', t_{r'}) = ND_\gamma(r', t)$, by correctness of P and as $t_{r'} = t < \infty$.) Since $r' \subseteq^t r$, any syncausal path in r' ending no later than at t is also a syncausal path in r , and therefore also $\tilde{e} \xrightarrow{\gamma_r} e$. Let $i \in I$. As $e \in ND_\gamma(r', t) = PND_\gamma^{r'}(i, t)$, we obtain $e \xrightarrow{\gamma_{r'}}(i, t)$. Note that any delivery $d \neq e$ along any syncausal path $e \xrightarrow{\gamma_{r'}}(i, t)$ satisfies both $\tilde{e} \xrightarrow{\gamma_r} d$ and $t_e < t_d$. Hence, by maximality of t_e , we have $d \notin ND_\gamma(r', t)$. Moreover, since $t_d \leq t$, we obtain $d \notin ND_\gamma(r')$ as well. Therefore, we obtain $e \xrightarrow{-\gamma_r}(i, t)$. \square

As noted in the previous section, the fact that a protocol is a finite-influence one need not dictate, in general, that the context is a bounded-syncausal-path one. Nonetheless, we now demonstrate a general technique that can be employed to show that many results regarding the existence of syncausal structures in bounded-syncausal-path contexts hold for finite-influence protocols in arbitrary contexts as well. Furthermore, many such results, which hold only for finite sets of agents in bounded-syncausal-path contexts, generalize, for finite-influence protocols, to hold for infinitely many agents as well.

Claim 9.14. *Corollary 9.4 (resp. Corollary 9.5) also holds when P is a finite-influence protocol, even when dropping the requirements for a bounded syncausal*

path from \tilde{e} to $(i, t_r(i))$ and for finiteness of J (resp. of each J_m). Under these conditions, the first part of that corollary guarantees, for arbitrary $i \in J$ (resp. $i \in J_1$) of our choosing, a finite* bound of the form $\tilde{b}(t_r(i), \sup(\hat{\delta}|_{J_2}), |E|)$ (resp. $\tilde{b}(t_r(i), n, (\sup(\hat{\delta}|_{J_m^2}))_{m=1}^n, (\inf(\hat{\delta}|_{J_m \times J_{m+1}}))_{m=1}^{n-1}, |E|)$), where E is a minimal $t_r(i)$ -retainable set that still guarantees i 's response at $t_r(i)$.

Proof. We present a proof of Claim 9.14 with regard to the first part of Corollary 9.4. The proofs with regard to the second part thereof, and to both parts of Corollary 9.5, are analogous.

Let $(\gamma, \tilde{e}, I, \delta)$ be a TCR-spec, Let $P \in TCR_\gamma\langle \tilde{e}, I, \delta \rangle$ be a finite-influence protocol, let $r \in R_\gamma^\tilde{e}(P)$ and let $J \subseteq I$ s.t. $\sup(\hat{\delta}|_{J_2}) < \infty$. Let $i \in J$ and let $E \in RND_\gamma^P(r, t)$ be minimal s.t. $t_{r \cap E}(i) = t_r(i)$. (Since P is a finite-influence protocol, there exists such a minimal E , and $|E| < \infty$.)

We first give a proof for the special case in which $|J| < \infty$, proving a weaker statement as we allow the bound to depend on $|J|$ for the time being.

Define $t \triangleq t_r(i)$ and $r' \triangleq r \cap^t E$. As $L_{r'}(\tilde{e}, (i, t_{r'}(i))) \leq |E| < \infty$, Corollary 9.4 may be applied to show the existence of an \tilde{e} -broom in r' by (no later than) $\tilde{b}(t_{r'}(i), \sup(\hat{\delta}|_{J_2}), |J|, L_{r'}(\tilde{e}, (i, t_{r'}(i)))) \leq \tilde{b}(t_r(i), \sup(\hat{\delta}|_{J_2}), |J|, |E|)$. As the broom event occurs no later than $t_{r'}(i) = t$ (and thus is in $ND_\gamma(r', t) = E$), this broom exists in r as well.

We now explain why the requirement for finiteness of $|J|$ may be dropped. Set $\tilde{b} \triangleq \tilde{b}(t_r(i), \sup(\hat{\delta}|_{J_2}), |E|, |E|) < \infty$ — the same bound as in the first part, substituting $|J|$ with $|E|$. If some $e \in E$ constitutes an \tilde{e} -broom for J by \tilde{b} , then we are done. Assume, by way of contradiction, that this is not the case. Thus, for each $e \in E$, there exists $j_e \in J$ s.t. $e \not\rightarrow^{\tilde{\gamma}} (j_e, \tilde{b})$. The set $\{j_e\}_{e \in E}$ is of size no greater than $|E| < \infty$, so the first part of this proof may be applied to it, yielding that there exists $e \in E$ that constitutes an \tilde{e} -broom for $\{j_e\}_{e \in E}$ by \tilde{b} — a contradiction. \square

As may be expected, Claim 9.14 yields generalized versions of the results of the corollaries that it generalizes.

Corollary 9.15. *In Corollaries 9.7, 9.8 and 9.9, the requirement for finiteness of agents may be relaxed to $\hat{\delta}$ being bounded from above on each strongly-connected*

* Once we allow J (resp. each J_m) to be infinite, though, the first part of that corollary requires the additional assumption that $\hat{\delta}$ is bounded from above on J (resp. on each J_m), for the guaranteed bound to be finite. (This assumption was redundant as long as J (resp. each J_m) was finite.)

component of G_δ , if SR_γ and TCR_γ are restricted to finite-influence protocols, and if the requirement of G_δ containing no negative cycles is generalized to δ being implementable.

9.3 Unbounded-Message-Delivery Contexts

We conclude this chapter by revisiting the scenario with which we opened it — that of the two generals who attempt to coordinate a simultaneous attack on a village. By now, they would probably be content with even an approximately simultaneous attack, so we consider this more generalized case. We use the insight this scenario has given us, which led us to define bounded-syncausal-path contexts, to prove the following impossibility result. This result both generalizes [11, Corollary 6.1.4], which shows that in a context that exhibits unbounded message delivery in a discrete-time model, common knowledge is unattainable, and strictly strengthens [11, Corollary 11.6.4], which shows that in a context that exhibits arbitrary message loss in a discrete-time model, ε -coordination based on an ND event is unattainable.

Corollary 9.16. *Let γ be a bounded-syncausal-path context satisfying $b_\gamma(i, j) = \infty$ for every $(i, j) \in N_\gamma^*$, let (I, δ) be an implementation-spec s.t. $I \subseteq \mathbb{I}_\gamma$, and let \tilde{e} be any ND event. If G_δ has any nontrivial (i.e. non-singleton) strongly-connected component, then $C_I^\delta(\odot^{\leq 0}(\tilde{e})) = (\emptyset)_{i \in I}$ in $\Omega_{R_\gamma(P)}$, for every $P \in \mathbb{P}_\gamma$.*

Proof. By revisiting the proof of Theorem 6.34, we may notice that it does not use all the properties of the timely-coordinated response problem. Let us define $TCR'_\gamma\langle\tilde{e}, I, \delta\rangle$, a strictly weaker[§] variant of the timely-coordinated response problem,[¶] as the set of all protocols $P \in \mathbb{P}_\gamma$ satisfying:

- In each $r \in R_\gamma^{\tilde{e}}(P)$, either no $i \in I$ responds, or they all do, each exactly once.
In each $r \in R_\gamma(P) \setminus R_\gamma^{\tilde{e}}(P)$, no $i \in I$ responds.
- In every run $r \in R_\gamma^{\tilde{e}}(P)$ in which all I respond, it holds that $t_r \in T(\delta)$.

It may be readily verified that the result of Theorem 6.34, using the exact same proof, still holds for TCR'_γ , for every triggered run during which all I respond.

* Such a context is said, in [11], to exhibit “unbounded message delivery”.

§ In the sense that $TCR_\gamma\langle\tilde{e}, I, \delta\rangle \subsetneq TCR'_\gamma\langle\tilde{e}, I, \delta\rangle$.

¶ While we have only defined the timely-coordinated response problem as based on external input events, it may be readily verified that all our results regarding it still hold if we allow it to be based on any ND event.

Furthermore, by Theorem 7.24 and by Corollary 7.31, the response logic defined in the latter is optimal in the sense that a full-information protocol endowed with it solves $TCR'_\gamma\langle\tilde{e}, I, \delta\rangle^*$ and moreover, for every full-information protocol $P \in TCR'_\gamma\langle\tilde{e}, I, \delta\rangle$, in each run thereof during which all I respond, replacing the response logic of P with this optimal response logic would still yield responses of all I in this run, and no response time would grow.

Assume, by way of contradiction, that there exists $P \in \mathbb{P}_\gamma$ and $i \in I$ s.t. $C_I^\delta(\otimes^{\leq 0}(\tilde{e}))_i \neq \emptyset$ in $\Omega_{R_\gamma(P)}$. Assume w.l.o.g. that P is a full-information protocol.[§] Furthermore, as we have not given any restrictions regarding the response logic of P , assume w.l.o.g. that P is endowed with the response logic defined in Corollary 7.31. By the above discussion, $P \in TCR'_\gamma\langle\tilde{e}, I, \delta\rangle$. Let $(r, t) \in C_I^\delta(\otimes^{\leq 0}(\tilde{e}))_i$. As in the proof of Theorem 7.24, $(r, t) \subseteq \diamond(K_j(C_I^\delta(\otimes^{\leq 0}(\tilde{e}))_j))$ for every $j \in I$, and thus all I respond in r according to P . Let $(j, k) \in I^2$ be a pair of distinct agents from the same strongly-connected component of G_δ . As $b_\gamma \equiv \infty$, the bound-guarantee relation is local-only.

Theorem 6.34, when applied to $\hat{\delta}$ and to paths alternating between j and k , implies, for every $n \in \mathbb{N}$, a syncausal path in r , from \tilde{e} , alternating, by $t_r(j)$ (which, by the above discussion, is finite), n times back and forth between j and k (as the bound-guarantee relation is local-only), which is impossible in a bounded-syncausal-path context — a contradiction. \square

Alternative proof ending. By Corollary 9.4,[¶] a broom for $\{j, k\}$ exists in r , contradicting the fact that the bound-guarantee relation is local-only. \square

Remark 9.17. *In Corollary 9.16, utilizing the technique employed in the proof of Claim 9.14, the requirement for a bounded-syncausal-path context may be dropped if P is restricted to be a finite-influence protocol.*

It is only fitting that Corollary 9.16 concludes the presentation of novel results in this work, as it demonstrates the added value of our dual approach to solving the timely-coordinated response problem, as the proof we have given thereto (regardless of the choice of ending) utilizes elements that are unique to each of the approaches we have taken.

* It should be noted that if G_δ is not strongly connected, then this is not generally true for the response logic defined in Corollary 6.35, nor is it generally true if we replace C_I^δ with \mathcal{C}_I^δ .

§ It may be readily verified that adding auxiliary variables to the state of each agent in order to turn P into a full-information protocol may only enlarge $C_I^\delta(\otimes^{\leq 0}(\tilde{e}))_i$.

¶ The same proof we have given when deducing Corollary 9.4 from Theorem 6.34 may be used to show that it, too, holds for TCR'_γ , for every triggered run during which all I respond.

Chapter 10

Deriving Previous Results

In this chapter, we show how some previously-known results may be derived from the novel results we have introduced in previous chapters.

10.1 General Ordered and Timed Responses

The problem of “general ordered” response was defined and studied by Ben-Zvi and Moses[5]. In this coordinated response problem, the relationships between the times of the responses of finitely-many agents is dictated by a given partial order relation \leq on classes of agents. The problems of ordered response, simultaneous response and ordered joint response, which were surveyed in the introduction to Chapter 6, are all special cases of this problem.

It may be readily seen that the general ordered response problem is a special case of the timely-coordinated response problem, for δ with the canonical form

$$(i, j) \mapsto \begin{cases} 0 & j \leq i \\ \infty & \text{otherwise.} \end{cases}$$

It should be noted that G_δ , sans the weights, is a DAG commonly used to describe the dual relation \geq in many applications. Hence, for δ with the above canonical form, every path $\bar{p} \in \mathcal{P}(G_\delta)$ is a weakly-decreasing tuple. For the above-mentioned special cases of general ordered response, this formulation coincides with their formulations as special cases of the timely-coordinated response problem, which we have given in the introduction to Chapter 6.

Ben-Zvi and Moses show that the syncausal structures underlying the general

ordered response problem (those that are guaranteed to exist in each triggered run of a solving protocol thereof, and that may be used to define an optimal response logic therefor) are centibrooms — one for each linearly-ordered chain of classes of agents, as guaranteed by Theorem 6.16 for such a chain in the ordered joint response problem. The second part of Corollary 9.5 reduces to this result in the special case of the general ordered response problem, and reduces to Theorem 6.16 in the special case of the joint ordered response problem.

Similarly, in the special case of the simultaneous response problem (resp. the tightly-timed response problem, with δ as in Remark 6.20), the second part of Corollary 9.4, together with Remark 6.36 and with Theorem 6.34, reduce to Theorem 6.12 (resp. Theorem 6.22).

Last but not least, we consider the case of ordered response (resp. weakly-timed response with δ as in Remark 6.18). In this case, the above-described partial order on I constitutes a linear ordering thereof. Therefore, all paths in G_δ are subpaths of the single strongly decreasing Hamiltonian path $\bar{p} = (p_m)_{m=1}^n \in \mathcal{P}(G_\delta)$. Therefore, for every $i \in I$, all paths $\bar{p}' \in \mathcal{P}(G_\delta)$ satisfying $p'_1 = i$ are subpaths of the suffix of \bar{p} that starts with i , which we denote by $\bar{p}_{\leq i}$. Thus, every corresponding (\bar{p}', δ) -traversing \tilde{e} -centipede is a subcentipede of a $(\bar{p}_{\leq i}, \delta)$ -traversing \tilde{e} -centipede by the time of i 's response. This path-traversing centipede is, in turn, simply an \tilde{e} -centipede by that time for $\bar{p}_{\leq i}^{rev}$. Thus, Theorem 6.34 and Corollary 6.35 reduce to Theorem 6.8 (resp. Theorem 6.21) in this case.

As previously noted, the proof that we presented to Theorem 9.12 readily generalizes to directly prove, among others, all the results surveyed in this section.

10.2 Common Knowledge and Variants

For the duration of the section, fix a context γ , a set of runs $R \subseteq \mathcal{R}_\gamma$, an event $\psi \in \mathcal{F}_R$ and a set of agents $I \subseteq \mathbb{I}_\gamma$. As noted above, while all previously-studied variants of common knowledge that are surveyed in the introduction to Chapter 7 are defined as fixed points of functions on \mathcal{F}_R , this is not the case with δ -common knowledge, which we define as a fixed point of a function on \mathcal{F}_R^I . Intuitively, as noted in that chapter, this stems from the asymmetry of δ -coordination with regard to the requirements posed on the various agents. Given this intuition, one may expect δ -common knowledge to reduce, for constant δ , to a non-tuple fixed point in some way. Indeed, if δ is a constant function, then it is straightforward to verify

that $K_i(C_I^\delta(\psi)) = K_i(\cap C_I^\delta(\psi))$ for every $i \in I$ and that $\cap C_I^\delta(\psi)$ is the greatest fixed point of $\cap f_\psi^\delta$. We now review the previously studied non-tuple variants of common knowledge and discuss when, and how, the above-described special case of δ -common knowledge for constant δ generalizes them.

When $\delta \equiv \infty$, then by definition, δ -coordination is equivalent to eventual coordination, $\cap f_\psi^\delta$ is the function presented in the first part of Theorem 7.14, and thus $\cap C_I^\delta(\psi) = C_I^\diamond(\psi)$. In addition, in this case Theorem 7.24 implies Theorem 7.14.

Reducing the results of δ -common knowledge to ε -common knowledge, for finite ε , is somewhat more delicate. Assume, for the remainder of this section, that $\delta \equiv \varepsilon$ for some finite $\varepsilon \geq 0$. (Recall that for $\varepsilon \equiv 0$, ε -coordination is equivalent to perfect coordination and Theorem 7.16 reduces to Theorem 7.10.)

In general, ε -coordination is a stricter condition than δ -coordination.* For a stable ensemble, though, δ -coordination is equivalent to ε -coordination. If we restrict ourselves to protocols exhibiting perfect recall, then by Corollary 8.8, the ensemble defined by δ -common knowledge is stable. If, in addition, ψ is stable, then it may be verified that the ensemble defined by ε -common knowledge is stable as well.† In this case, by Lemma 8.12, $C_I^\delta(\psi)$ is the greatest fixed point of g_ψ^δ and thus, $\cap C_I^\delta(\psi)$ is the greatest fixed point of $\cap g_\psi^\delta$. Analogously to the proof of Lemma 8.12, but in a less cumbersome way (as $\delta < \infty$), it may be shown that in this case $C_I^\varepsilon(\psi)$ is the

* This stems from two main “reasons”:

1. δ -coordination is defined using $\odot^{\leq \delta(i,j)}$ rather than $\odot^{[-\delta(j,i), \delta(i,j)]}$, which we define to mean “at some time no earlier than $-\delta(j,i)$ from now and no later than $\delta(i,j)$ from now”. It may be readily seen that all the results in this work hold for such a definition as well, as long as this replacement is performed in the definition of f_ψ^δ as well. The only difference is that Claim 8.7, stating that δ -common knowledge is stable, requires also stability of ψ and perfect recall in this case, and is proven by showing that $\odot^{\leq 0}(C_I^\delta(\psi)) \leq f_\psi^\delta(\odot^{\leq 0}(C_I^\delta(\psi)))$ and by applying the second part of Lemma 7.23.
2. δ -coordination is based on pairwise constraints. The results presented in this work may be quite readily generalized to deal with arbitrary timing constraints of various natures, such as, e.g. for some $J \subseteq I$, “For every $i \in J$ and for every $(r, t) \in \mathbf{e}_i$, there exists a time interval $T \subseteq \mathbb{T}$ of length at most δ_J , s.t. $t \in T$ and s.t. there exist $(t_j)_{j \in J} \in T^J$ satisfying $(j, t_j) \in T$ for every $j \in J$ ”. (Whatever the timing constraints are, the generalized definition of $f_{\psi_i}^\delta$ simply intersects on all constraints pertaining to i .) Under such a generalization, ε -coordination is equivalent to δ -coordination, when setting $\delta_I \equiv \varepsilon$ in the above constraint example, and when providing no further constraints. Furthermore, in this case the generalization of f_ψ^δ satisfies that $\cap f_\psi^\delta$ is the function presented in the first part of Theorem 7.16, and thus the appropriate generalization of Theorem 7.24 reduces to Theorem 7.16.

† The key observation required for showing this is that $\odot^{\leq 0}(C_I^\varepsilon(\psi)) \subseteq E_I^\varepsilon(\psi \cap \odot^{\leq 0}(C_I^\varepsilon(\psi)))$, given stability of ψ and perfect recall.

greatest fixed point of $\cap g_\psi^\delta$ as well, and thus $\cap C_I^\delta(\psi) = C_I^\varepsilon(\psi)$.^{*}

In the absence of stability of ψ , or in the absence of perfect recall (at least of the “relevant events”), things stop working so well. Indeed, as noted above, in such cases δ -coordination does not necessarily coincide with ε -coordination, and consequently, examples may be constructed in which the ensembles defined by ε -common knowledge and by δ -common knowledge differ.

^{*} Another way to derive this equality is by using [11, Exercise 11.17(d)], which shows that, for every $i \in I$, when ψ is stable and given perfect recall, $K_i(C_I^\varepsilon(\psi)) = K_i(\cap_{n \in \mathbb{N}} (\odot^\varepsilon E_I)^n(\psi))$, to which (8.2) reduces when $\delta \equiv \varepsilon$. It should be noted, though, that the proof hinted to by [11, Exercise 11.17(d)] strongly relies on a discrete modeling of time, and breaks down in a continuous-time model, unlike the proof that we sketch above.

Chapter 11

Discussion and Open Questions

11.1 A Qualitative Comparison of Approaches

Throughout this work, reasoning alternated between two approaches, which are based on different motivations and thus were previously studied only separately. In Chapter 8, though, we showed that despite the vast conceptual difference between these two approaches, they in fact yield equivalent results for the timely-coordinated response problem. Nonetheless, this conceptual gap makes each approach convenient for different purposes. Consequently, we have utilized each approach to attack a different set of problems in Chapter 9. Moreover, in Corollary 9.16 we concurrently harnessed both approaches to obtain a strengthened version of a previously-known result.

The strength of the syncausal approach, similarly to that of Lamport’s asynchronous causality[17] that it generalizes, is its constructiveness and concreteness. These properties make it ideal for graphical visualization of runs and for algorithm design. However, their price is the need to adapt and specifically tailor the general results for each model flavour, as we have done in Chapter 9.

Conversely, the strength of the fixed-point approach lies in its generality and in its high level of reasoning. These properties make proofs that follow this approach far less cumbersome, and far more general, due to the fact that, as we have seen, the concept of knowledge effectively hides the minute details of the model in question. The downside of this is the fairly large gap, both between a fixed-point definition and a constructive definition, as we have seen in Chapter 8, and moreover — between a constructive knowledge-based definition and concrete implementation, as studied by Ben-Zvi and Moses[6, 5, 7, 8] following Chandy and Misra[9].

Given these observations, it is not surprising that the results we obtained using fixed-point analysis in Chapter 7 are far more general (and can even be further generalized, as noted in Chapter 10), than the results we obtained using syncausal analysis in Chapter 6. However, it is the latter that were easier for us to turn into a concrete algorithm in Example 9.6, and into a concrete condition, in terms of required guarantees on message delivery times, for solvability of the timely-coordinated response problem in a given context in Corollaries 9.9 and 9.15.

11.2 On Generalizations

Much of this work is based on two generalizations of known approaches. Nontrivial generalizations tend to have a sneaky property: on one hand, a conceptual leap is required in order to achieve them, while on the other hand, once they are achieved, this leap, in hindsight, seems almost obvious.

The work of Ben-Zvi and Moses on syncausal analysis[6, 5, 7, 8] is implicitly intertwined with an insight, which holds for all of the problems they define and analyze: Each of these problems has one, succinctly-describable*, syncausal structure underlying it.[§] Indeed, when we started looking at simple two-agent cases of what would eventually become the timely-coordinated response problem, we attempted to find such a simple structure, or possibly only a few simple structures.

Moreover, the asymmetry inherent in the syncausality and bound guarantee relations expresses itself in the problems defined by Ben-Zvi and Moses, in that the timing dependency between the response times of two agents in these problems may only be single-sided, which allows the response of one of these agents to not depend on the response of the other.[¶] Indeed, as noted in Chapter 6, this is the main difference, both conceptually and technically, between the response problems defined and studied by Ben-Zvi and Moses[6, 5, 7, 8] and the timely-coordinated response problem, which we have defined and analyzed in this work. As we have seen, it

* One may almost claim that the description should be linear in the number of agents.

[§] The one exception to this is the general ordered response problem, which is treated by Ben-Zvi and Moses[5] as a conjunction of *independent* ordered joint response problems, and thus their solution consists of a conjunction of the syncausal structures (i.e. the centibrooms) underlying each of these ordered joint response problems.

[¶] Actually, their analysis, as we have seen, also allows a precise timing dependency between two response times (i.e. a specified fixed time difference), which allows them to be treated, in a sense, as a single response in the solution of the problem. Their analysis does not, however, allow a mutual (i.e. double-sided) non-precise dependency.

is the presence of mutual non-precise dependencies, that changes the “rules of the game” from revolving around one, fairly simple, syncausal structure to revolving around infinitely many, or alternatively finitely many yet very complex*, syncausal structures, which may not, in general, be replaced by simpler or fewer structures.[§]

As we commented earlier, up until now the fixed-point approach has only been applied to problems whose description exhibits an inherent symmetry between the agents, in the sense that it is invariant under permutations on the set of agents. As noted in Chapter 7, it is the absence of this symmetry that “twisted our arms” and conceptually necessitated the nontrivial jump from searching for a fixed point of a scalar function to searching for a fixed point of a vectorial function. Moreover, even after the realization that this is the way to go, this vectorial treatment was the main technical obstacle in our fixed-point analysis.

11.3 Open Questions and Further Directions

Throughout this work, we assume a context in which the eventual delivery of any message is guaranteed. In many models that do not present this behaviour, arbitrarily long syncausal paths present themselves with zero (or very small) probability, effectively displaying a behaviour similar to that of the class of bounded-syncausal-path contexts, which we defined in Chapter 9. It may be interesting, therefore, to develop such probabilistic models and to check whether the results given in Chapter 9 for bounded-syncausal-path contexts may be applied to such models, if only to yield either probabilistic results or impossibility results.

Corollary 9.7 implies that in a shared-clock, bounded-syncausal-path context, solving an “almost-simultaneous” response problem is not any more possible than solving a simultaneous response problem. Moreover, if both are solvable, then (9.1) implies that in the worst-case scenario, the time of the latest of the responses in the optimal solutions to both problems is the same.[¶] As was noted in Corollary 9.15,

* In fact, arbitrarily complex even for two agents.

§ Such a “replacement” is performed e.g. in the proof of Ben-Zvi and Moses[6, 5] for the first part of Theorem 6.12. In this work, we obtain such replacements for some special cases of the timely-coordinated response problem in Corollaries 9.4 and 9.5 and in Claim 9.14.

¶ By Corollary 9.8, similarly relaxing the tight constraints of an ordered joint, or tightly-timed, response problem also does not make it solvable in any additional contexts, nor does it improve the worst-case time of the latest of the responses.

|| Nonetheless, in all cases the relaxed version may be solved significantly faster than the original one in many runs, as illustrated in Figure 5. (Thus, ACME’s engineers were on the right track in Example 1.1.) It would be interesting to give this observation a precise meaning in a probabilistic model, perhaps in terms of average-case response time.

even without the assumption of a bounded-syncausal-path context, these results hold as long as we make some reasonable assumptions regarding finiteness of memory or of processing power of each agent. Furthermore, Corollary 9.16 shows that in the lack of any delivery guarantees, none of these problems are solvable under such reasonable assumptions. Nonetheless, it has been shown in [11, Subsection 11.2.1] that in models without a shared clock, “up-to- ε ” coordination may be possible even when perfect coordination is not. It would be interesting to see whether the machinery presented in this work may be applied, perhaps in some extended or generalized form, to shed new light on models in which the clock is not shared.

In Chapter 7, we defined and analyzed δ -coordination, a generalization of several forms of coordination defined and analyzed by Halpern and Moses[14] and by Fagin et al.[11, Section 11.6]. In Chapter 10, we noted that some special cases of ε -coordination (another form of coordination defined in [14, 11]) are not generalized by δ -coordination. While, as we noted there, our definition of δ -coordination, along with all our results regarding δ -common knowledge, may be quite readily generalized to deal with additional forms of coordination constraints, including those of ε -coordination, it remains to be seen whether such generalizations are of any real added value. In this context, it is worth to recall the difficulty we encountered in Chapter 7, in giving a succinct characterisation to the ensemble defined by δ -common knowledge (or by ε -common knowledge, for that matter) of an event along the lines of “the greatest δ/ε -coordinated ensemble satisfying...”. This difficulty, coupled with slight differences in the properties of δ/ε -common knowledge, raises the following question: have we truly given the “right”, “from the book” definition for δ -common knowledge? (Similarly, have Halpern and Moses[14], and Fagin et al.[11, Section 11.6], given the “right” one for ε -common knowledge?) or is a similar, yet succinctly characterisable, fixed-point definition still waiting to be phrased?

We conclude this work with a comment about fixed points. As we have seen, fixed-point analysis of coordination is useful in a significantly broader range of cases than previously thought. Many systems around us, from subatomic physical systems to astrophysical ones, and from animal societies to some stock markets, exist in some form of equilibrium fixed point, possibly reached as a result of a long-forgotten spontaneous symmetry breaking. This leads us to conjecture that describing distributed algorithms as fixed points may potentially be of much further advantage and provide us with additional insights that are yet to be discovered.

Bibliography

- [1] E. A. Akkoyunlu, K. Ekanadham, and R. V. Huber. Some constraints and tradeoffs in the design of network communications. In *Proceedings of the 5th ACM Symposium on Operating Systems Principles (SOSP '75)*, pages 67–74, 1975.
- [2] A. Arnold and D. Niwiński. *Rudiments of μ -Calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, Netherlands, 2001.
- [3] R. J. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3):263–300, 1999.
- [4] J. Barwise. Three views of common knowledge. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK '88)*, pages 365–379, 1988.
- [5] I. Ben-Zvi. *Causality, Knowledge and Coordination in Distributed Systems*. PhD thesis, Technion, Israel Institute of Technology, Haifa, Israel, 2011.
- [6] I. Ben-Zvi and Y. Moses. Beyond Lamport’s happened-before: On the role of time bounds in synchronous systems. In *Proceedings of the 24th International Symposium on Distributed Computing (DISC 2010)*, pages 421–436, 2010.
- [7] I. Ben-Zvi and Y. Moses. On interactive knowledge with bounded communication. *Journal of Applied Non-Classical Logics*, 21(3-4):323–354, 2011.
- [8] I. Ben-Zvi and Y. Moses. Agent-time epistemics and coordination. Unpublished manuscript, available on arXiv, 2012.
- [9] K. M. Chandy and J. Misra. How processes learn. *Distributed Computing*, 1(1):40–52, 1986.

- [10] B. A. Coan, D. Dolev, C. Dwork, and L. Stockmeyer. The distributed firing squad problem. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing (STOC '85)*, pages 335–345, 1985.
- [11] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. The MIT Press, Cambridge, MA, USA, 1995.
- [12] M. F. Friedell. On the structure of shared awareness. *Behavioral Science*, 14(1):28–39, 1969.
- [13] J. Gray. Notes on data base operating systems. In R. Bayer, R. M. Graham, and G. Seegmüller, editors, *Operating Systems, An Advanced Course*, volume 60 of *Lecture Notes in Computer Science*, pages 393–481. Springer-Verlag, Berlin, Germany / Heidelberg, Germany / New York, NY, USA, 1978.
- [14] J. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
- [15] S. C. Kleene. *Introduction to Metamathematics*. North-Holland Publishing Company, Amsterdam, Netherlands, 1952.
- [16] I. I. Kolodner. On completeness of partially ordered sets and fixpoint theorems for isotone mappings. *American Mathematical Monthly*, 75(1):48–49, 1968.
- [17] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.
- [18] L. Lamport. Using time instead of timeout for fault-tolerant distributed systems. *ACM Transactions on Programming Languages and Systems*, 6(2):254–280, 1984.
- [19] D. Lewis. *Convention, A Philosophical Study*. Harvard University Press, Cambridge, MA, USA, 1969.
- [20] R. Michel. A categorical approach to distributed systems, expressibility and knowledge. In *Proceedings of the 8th ACM Symposium on Principles of Distributed Computing (PODC '89)*, pages 129–143, 1989.
- [21] E. F. Moore. The firing squad synchronization problem. In E. F. Moore, editor, *Sequential Machines: Selected Papers*, pages 213–214. Addison-Wesley, Reading, MA, USA, 1964.
- [22] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.

Appendices

Appendix A

A Continuous-Time Model

In this appendix, we describe a novel continuous-time model for which the results in this work hold verbatim. In order to avoid repetitions, we only describe the differences between this model and the model presented in Chapter 3.

A.1 Context Parameters

In the continuous-time model, as in the discrete-time one, we denote a context by a tuple $\gamma = (\mathcal{G}_\gamma = (\mathbb{I}_\gamma, N_\gamma, b_\gamma), (S_i)_{i \in \mathbb{I}_\gamma}, \tilde{E}_\gamma, (i_{\tilde{e}})_{\tilde{e} \in \tilde{E}_\gamma})$. In this case, though, \mathcal{G}_γ is a weighted directed graph with positive, real or infinite, weights. Additionally, we define the set of times as $\mathbb{T} \triangleq \mathbb{R}_{\geq 0}$. Our reasons for this somewhat unorthodox approach to modeling time (in a continuous fashion) hopefully become apparent throughout Chapter 9.

A.2 Timers

For the duration of this section, fix an agent $i \in \mathbb{I}_\gamma$. As time is continuous in our model, we wish to define when i is allowed to act. We say that i is “enabled” (to act) at $t \in \mathbb{T}$ if either $t = 0$ or t is a supremum of a set of times, at each of which i observed an event. (Intuitively, this means that either i observed an event at exactly t , or i observed infinitely many events, whose respective times converge to t in a monotonically-increasing fashion.)

In order to allow an agent $i \in I$ to make sure it has a chance to act at a certain time, we introduce a new type of action, and a new type of event. At any time t at which i is enabled, i may set a timer for a future time $t' \in \mathbb{T}$ s.t. $t' > t$. If i

sets such a timer, and if i is not enabled at any time between t and t' , then a timer ring event is observed by i at t' , thus enabling it at t' . It should be noted that it is also possible, though somewhat less intuitive, to not introduce timers, but rather to enable every agent at every time.

Due to the introduction of timers, the set of all possible states of the environment becomes $S_e \triangleq 2^{\tilde{E}_\gamma} \times 2^{\mathcal{M} \times \mathbb{T} \times N_\gamma} \times 2^{\mathbb{I}_\gamma}$ (the last element is a subset of the agents, for which timers ring at the current time), and the set of possible actions which may be taken by i at any time at which it is enabled becomes $A_i \triangleq S_i \times 2^{\mathcal{M} \times \{j \in I \mid (i,j) \in N_\gamma\}} \times \{t' \in T \mid t' > t\} \times \{\text{false}, \text{true}\}$.

A.3 Agent States

In order to define the state of an agent “just before” a time $t \in \mathbb{T}$, we assume, for each $i \in \mathbb{I}_\gamma$, the existence of a pseudo-limit function $\lim_i : S_i^{(-1,0)} \rightarrow S_i$ satisfying:

1. $\lim_i(f) = \lim_i(g)$, for every $f, g : (-1, 0) \rightarrow S_i$ s.t. $f|_{(-\varepsilon, 0)} = g|_{(-\varepsilon, 0)}$ for some $\varepsilon > 0$.
2. $\lim_i(f) = s$, for a constant function $f \equiv s \in S_i$.

There are a number of other natural properties which one may expect from \lim_i (e.g. invariance to composition with monotone continuous functions from $(-1, 0)$ to itself, which have 0 as their limit as 0), but we do not require any such properties for the results of this work.

For a full-information protocol, in which the state of each agent $i \in \mathbb{I}_\gamma$ at any time $t \in \mathbb{T}$ uniquely determines the full details of every event observed by i up until, and including, t , a natural pseudo-limit function is (infinite) union of sets of events. In general, functions such as union, logical or, max, and min, are useful building blocks for pseudo-limit functions for many intuitive protocols.

We define full-information protocols in this model in a similar way to Chapter 3, although in this model an agent only sends out messages in a full-information protocol when it is enabled. Two full-information protocols may thus differ not only in their response logics, but also in their timer-setting logics. Thus, there does not necessarily always exist an isomorphism between the sets of runs of two full-information protocols that preserves the set of ND events. Nonetheless, it is still true that given a protocol P , there exists a full-information protocol P' , s.t. there

is a natural monomorphism from $R_\gamma(P)$ into $R_\gamma(P')$, which preserves both the set of ND events, and all responses.

A.4 Runs

We are now ready to redefine the properties that a function $r : \mathbb{T} \rightarrow S_e \times \prod_{i \in I} S_i$ must satisfy in order to constitute a legal run of a protocol $P = ((\tilde{S}_i, P_i))_{i \in I} \in \mathbb{P}_\gamma$:

- Agent state consistency with local protocol: Let $i \in I$ and $t \in \mathbb{T}$. If $t > 0$, set $s_i = \lim_i(r_i(t + \cdot))$. (This is well defined even when $t < 1$, because \lim_i only depends on the values of its argument in a left neighbourhood of 0.) Intuitively, s_i is the state of i “just before” t .^{*} If $t = 0$, then s_i may be any of the initial states \tilde{S}_i .

If i is not enabled at t (this condition depends only on the environment states before and at t), then $r_i(t) = s_i$ must hold.

If i is enabled at t , then $r_i(t)$ must equal the first part of the output of P_i , when evaluated on s_i and on the events observed by i at t . (Once again, the other parts thereof determine the actions of i at t .)

- Environment state properties:
 1. The requirements regarding external inputs and message deliveries are unchanged.
 2. Timer events: A timer event for $i \in I$ occurs at time $t \in \mathbb{T}$ iff there exists $t' \in \mathbb{T}$ such that $t' < t$ and such that i set, at t' , a timer to ring at t , and i was not enabled during (t, t') .

A.5 Excluding Degeneracies

While, in the discrete-time model, a given “partial run” $r : \{t \in \mathbb{T} \mid t \leq 1\} \rightarrow S_e \times \prod_{i \in I} S_i$ of a protocol $P \in \mathbb{P}_\gamma$ may be inductively “rolled forward” to create a full (infinite) run (see, e.g. Claim 6.27), this may no longer be the case in a continuous-time model with infinitely many agents. Intuitively, consider such a partial run, in which infinitely many messages $\{m_k\}_{k=1}^\infty$ are sent to some agent $i \in \mathbb{I}_\gamma$ before time 1,

^{*} Note that if $r_i|_{(t', t)}$ is constant for some $t' \in \mathbb{T}$ s.t. $t' < t$, then s_i equals this constant value.

but are not yet delivered by that time. Assume, furthermore, that for every $k \in \mathbb{N}$, the message m_k is guaranteed to be delivered no later than at $1 + \frac{1}{k}$. It is not clear how to “roll the run forward”, even for a fraction of a time unit. Similarly, if the delivery guarantee for each m_k is at $2 + \frac{2}{k}$, then it is not clear that it is possible to roll the run forward while avoiding early deliveries, as required in Claim 6.27. Indeed, if it is not possible to do so, then in any run that is indistinguishable from this partial run up to time 1, it is *deterministic* that some early delivery takes place between times 1 and 2, effectively voiding the non-determinism of some early deliveries, possibly allowing them to be predicted before they occur.

In order to avoid degeneracies such as those described above, and thus sufficiently maintain the non-deterministic nature of the events that we call “ND events”, we axiomatically make the following assumptions regarding the richness of the set of runs of any protocol $P \in \mathbb{P}_\gamma^*$

- No foresight: For every $r \in R_\gamma(P)$ and for every $t, d \in \mathbb{T}$, there exists a run $r' \in R_\gamma(P)$, satisfying:
 1. $r'|_{[0,t]} = r|_{[0,t]}$.
 2. No external inputs are triggered in r' after t .
 3. Any message delivered *early* in r' after t is delivered no less than d time units after it is sent.
- No extrasensory perception: For every $r \in R_\gamma(P)$, for every $t \in \mathbb{T}$ and for every $i \in \mathbb{I}_\gamma$, it holds that $PND_\gamma^r(i, t) \in RND_\gamma^P(r, t)$.[§]

The two above assumptions, regarding non-determinism of future events and independence of past events, respectively, imply that agents may not predict the occurrence of certain ND events. These assumptions suffice for most of our analysis. For some arguments, though, we require some guarantee that agents may not predict the *absence* of ND events. The following assumption, regarding non-determinism and independence of present events, provides such a guarantee and, if $R_\gamma(P) \neq \emptyset$, complements the above assumptions in the strongest way possible in some sense.

* As is discussed in Section 6.2, both of these assumptions may be shown to hold for the discrete-time model presented in Chapter 3. Furthermore, if $\inf(b_\gamma) > 0$, then an inductive argument, rolling a run forward $\inf(b_\gamma)$ time units at a time, may be used to show that these assumptions also hold for protocols implemented using certain “nice” pseudo-limit functions such as the union pseudo-limit function described above, and for arbitrary protocols if $|\mathbb{I}_\gamma| < \infty$.

[§] See Section 6.2 for the definitions of RND and of PND . (We allow ourselves to state this assumption in terms of syncausality, as we only utilize it in our syncausal analysis.)

- Eternal vigilance: For every run $r \in R_\gamma(P)$, for every $t \in \mathbb{T}$, for every set $E \subseteq \tilde{E}_\gamma$ of external inputs that are not triggered in r before t and for every set of M of potential early deliveries at t in r (i.e. messages sent before t , not delivered before t , and with a delivery guarantee greater than t), there exists a run $r' \in R_\gamma(P)$, satisfying:

1. $r'|_{[0,t)} = r|_{[0,t)}$.
2. $r'_e(t) = (E, M)$.

Although it may be readily verified that the above assumption holds in the discrete-time model presented in Chapter 3, this assumption is restrictive for a continuous-time model, as e.g. it does not hold for some naturally-occurring models, such as models with minimum bounds on delivery times. Moreover, this assumption also hinders the possibility of capturing a discrete-time model using our continuous-time model. (See the next section for more details.) For these reasons, we replace this assumption with the following, weaker assumption, which stems from combining the “eternal vigilance” and “no foresight” assumptions when $R_\gamma(P) \neq \emptyset$:

- Stand-alone external inputs: For every external input $\tilde{e} \in \tilde{E}_\gamma$, there exists a run $r \in R_\gamma^\tilde{e}(P)$, in which no ND events other than \tilde{e} occur before or at $t_{\tilde{e}}$.

A.6 Modeling Discrete Time

Now that we finished describing this model, it should be noted that discrete-time models, such as the one presented in Chapter 3, may be captured by this model. As an example, an integral-time model may be modeled by setting $b_\gamma(i, j)$ to integral (or infinite) values for every $(i, j) \in N_\gamma$, by forcing the environment to perform ND events only at integral times (i.e. removing from $R_\gamma(P)$ any runs in which any ND events occur at non-integral times), and by allowing timers to be set for integral times only. (Or, alternatively, by dropping timers altogether, and enabling every agent at every integral time.) We consider environment constraints, such as “all ND events occur at integral times”, or “any message may be delivered, at the earliest, ε after it was sent”, as integral parts of the model (just as the delivery bounds are). Care should be taken to make sure that such constraints do not interfere with the assumptions of the previous section.

תקציר

באלגוריתם מבוזר, מעבדים (סוכנים) מרובים פועלים להגשמת מטרה משותפת. לעתים קרובות, פעולות של סוכנים מסוימים תלויות בסיום הביצוע (אם לא גם בתוצאה) של פעולות של סוכנים אחרים. התלויות ההדדיות בין תזמוני הפעולות של סוכנים שונים מניעות מחקר של שיטות לתזמון מתואם של הפעולות הללו.

בעבודה זו, אנו מגדירים ומנתחים באופן מתמטי בעיה חדשנית של תאום בין סוכנים מרובים, שאותה נכנה "תזמון תגובה מתואם". בבעיה זו, הפרש הזמנים בין כל זוג פעולות ניתן לחסימה מלעיל ו/או מלרע. בעיה זו מכלילה מספר בעיות תאום קלאסיות שהוגדרו ונחקרו ע"י הלפרן ומוזס ומספר בעיות תאום שהוגדרו ונחקרו לאחרונה ע"י בן-צבי ומוזס.

אנו פותרים את בעית תזמון התגובה המתואם באופן מיטבי (כלומר, מגדירים פרוטוקול פותר מיטבי עבורה) בשתי דרכים: באמצעות הכללה של גישת נקודת השבת של הלפרן ומוזס ובאמצעות הכללה של גישת הנסיבתיות הסנכרונית של בן-צבי ומוזס. יתרה מכך, אנו מראים באופן מפורש את השקילות בין הפתרונות הנובעים משתי הגישות הללו, על אף ההבדלים התפיסתיים הנרחבים ביניהן. באמצעות שילוב של שתי הגישות, אנו מקישים גרסאות מחוזקות של תוצאות שהיו ידועות עבור כמה מקרים פרטיים של הבעיה שהוגדרו בעבר.

הניתוח שלנו מתבצע תחת הנחות מועטות: אנו עובדים במודל שהזמן בו רציף ושכמות הסוכנים בו יכולה להיות אינסופית. התוצאות הכלליות שאנו מקבלים עבור המודל הזה מצטמצמות לכדי תוצאות חזקות יותר עבור מודלים אשר בהם הזמן בדיד וכמות הסוכנים סופית. על מנת לזקק את התכונות של המודלים הללו שהנן משמעותיות לצורך צמצום זה, אנו מגדירים מספר מחלקות חדשניות של מודלים המופיעים באופן טבעי (כולן מכלילות מודלים עם זמן בדיד וכמות סופית של סוכנים), אשר במובן מסוים מפרידות בין התוצאות השונות. אנו חוקרים את בעית תזמון התגובה המתואם תחת המודלים הללו ומציגים בהם הן פתרון מיטבי מעשי יותר עבור הבעיה והן תנאי מפתיע בפשטותו לפתירותה.

לסיום עבודה זו, אנו מראים כיצד התוצאות שלנו עבור בעית תזמון התגובה המתואם מכלילות את התוצאות שהיו ידועות עבור המקרים הפרטיים של הבעיה שנחקרו בעבר. כמו כן, אנו מציגים מספר שאלות פתוחות וכווני מחקר עתידיים.

שלמי תודה

ברצוני להודות למנחה שלי, פרופ' גיל קלעי, על ידידותו ועל תמיכתו לאורך לימודיי. ברצוני להודות למנחה השני שלי, פרופ' יורם מוזס, על ידידותו, על שהכיר לי את הנושא המרתק של ידע ועל סבלנותו בעודי מפגיז את תיבת הדוא"ל שלו בהודעות במהלך השעות הקטנות של הלילה.

בנוסף, ברצוני להודות לחבריי, למשפחתי ואחרונה חביבה, לבת זוגי עלי, על תמיכתם, על עידודם ועל חברותם, ופשוט על היותם שם.

להורי, שלמדוני שידע הוא הדבר האחד שאיש לעולם לא יוכל
לקחת ממני,

לאחיי, שידעים הכל,

ויותר מכל, לעלי, האירוע היציב של חיי; נקודת השבת של חיי.

(עבודה זו ככל הנראה לעולם לא הייתה נכתבת אלמלא ההרפתקה שלה בקיוטו.)

האוניברסיטה העברית בירושלים
הפקולטה למתמטיקה ולמדעי הטבע
מכון אינשטיין למתמטיקה

תזמון מתואם במערכת מרובת סוכנים

עבודת גמר מחקרית לשם קבלת תואר מוסמך במדעי הטבע

מאת:

ינאי א. גונצ'רובסקי

מנחים:

פרופ' יורם מוזס

הפקולטה להנדסת חשמל
הטכניון, מכון טכנולוגי לישראל

פרופ' גיל קלעי

אדר התשע"ב