Catch-up Scheduling for Childhood Vaccination

Faramroze G. Engineer
School of Mathematical and Physical Sciences
University of Newcastle
Australia
faramroze.engineer@newcastle.edu.au

Abstract

In this report, we outline the development of the core optimization technology used within a decision support tool to help providers and caretakers in constructing catch-up schedules for childhood immunization. These schedules ensure that a child continues to receive timely coverage against vaccine preventable diseases in the likely event that one or more doses have been delayed. We develop a Dynamic Programming algorithm that exploits the typical size and structure of the problem to construct optimized schedules at almost the click of a button. In using an optimization based algorithm, our approach is unique not only in methodology but also in the information, strategy and advice we can offer to the user.

The tool is being advocated by both the CDC and the American Academy of Pediatrics (AAP) as a means of encouraging caretakers and providers to take a more proactive role in ensuring timely vaccination coverage for children, as well as ensuring the accuracy and quality of a catch-up regime.

1 Introduction

With the goal of ensuring timely and accurate administration of vaccines, the Advisory Committee on Immunization Practices (ACIP) of the CDC together with the AAP and the American Academy of Family Physicians (AAFP) annually publish a *recommended* immunization schedule for children aged 0 to 6 years (see (CDC 2008)). For a child who misses the recommended time for a dose, a healthcare professional faces the challenging task of constructing a *catch-up* schedule for that child under certain rules and guidelines for the administration of the remaining doses. These rules and guidelines specify the feasible number, timing and spacing of doses of each vaccine based on the child's age, the number of doses already received, and the child's age when each dose was previously administered (see (CDC 2008) for a summary of guidelines for catch-up immunization).

Immunization programs have a significant impact on public health and have been shown to be one of the most beneficial and cost effective disease prevention measures ((Zhou et al. 2005) and (Maciosek et al. 2006)). Although the majority of school going children in the United States that are six years and over are deemed covered against vaccine preventable diseases, most do not receive the optimal protection due to incomplete, untimely or erroneous vaccination. A comprehensive study carried out by (Luman et al. 2002) found that only 9% of children surveyed received all of their vaccinations at the recommended times and that only half received all their recommended doses by their second birthday. data gathered as part ofThe introduction of new vaccines to the recommended schedule adds complexity and the potential for deterioration in the overall timeliness of vaccination. Once a child falls behind the recommended schedule, statistics indicate that they often do not catch-up until close to reaching a school going age when an accelerated regime is most likely administered to meet the minimum coverage mandated by most schools.

Several factors contribute to poor and untimely vaccination rates. Some, such as parental misunderstanding and logistical difficulties affected by various environmental and socioeconomic factors are generally difficult to address and remedy. However, the problem is often exacerbated by incomplete and inaccurate catch-up schedules constructed by healthcare professionals. Constructing an accurate catch-up schedule is both a challenging and time consuming task. It therefore comes as no surprise that healthcare professionals struggle to manually construct catch-up schedules that reflect the best possible coverage for a child ((Cohen et al. 2003) and (Irigoyen et al. 2003)) and providers often fail to identify opportunities to vaccinate a child who may be at a clinic for purposes other than vaccination ((Holt et al. 1996) and (Szilagyi et al. 1993)). The complexity of the task is highlighted by the survey carried out by (Cohen et al. 2003) in which healthcare professionals were asked to construct catch-up schedules for 6 different hypothetical scenarios describing children who have fallen behind. On average, only 1.83 out of the schedules constructed for the 6 scenarios were deemed correct.

In this report, we investigate the catch-up scheduling problem and outline a Dynamic Programming (DP) algorithm that has been successfully adopted within a tool (downloadable from `www.cdc.gov/vaccines/recs/scheduler/catchup.htm`) developed jointly by CDC and Georgia Institute of Technology to help caretakers and providers make timely and accurate decisions with regards to childhood vaccination.

In what follows, we give a precise description of the catch-up scheduling problem in §2, outline the dynamic programming algorithm in §3, present solutions obtained for two real-life scenarios in §4, and relay some preliminary statistics about the use of the tool in practice and initial feedback from both physicians and parents in §5.

2 Problem Description and Notation

Given the current age of a child and their vaccination history (i.e., the number and timing of doses of each vaccine already administered), the catch-up scheduling problem is one of constructing a schedule for the remaining doses so that each dose is scheduled within the minimum and maximum age for that vaccine and dose, and the time separation between (not necessarily successive) doses of the same vaccine does not violate a certain minimum gap. This minimum gap may vary by vaccine, dose, current age and/or age at which some previous dose is administered. For example, the minimum gap between the second and third dose of *Hib* is 4 weeks if the current age is < 12 months, and the minimum gap is 8 weeks if the current age is ≥ 12 months and the second dose is administered at age < 15 months. In addition to

Table 1: The spacing between the first dose and remaining doses of *PCV*

dose i	dose $j > i$	age i is administered	age j is considered for administration	min gap between i and j
1	2	$<$ 12 months	any age	4 weeks
1	3	$<$ 12 months	any age	8 weeks
1	4	$<$ 12 months	any age	16 weeks
1	2	$<$ 24 but $\geq$ 12 months	[24, 60 months)	8 weeks
1	3,4	$\geq$ 12 months	[24, 60 months)	∞
1	2, 3, 4	$\geq$ 24 months	any age	∞

regulating the gap between doses of the same vaccine, doses of *live*[1] vaccines can only be administered during the same visit or at least a certain number of fixed days apart (28 days under the current guidelines). Finally, the number of *simultaneous administrations* (i.e., number of vaccinations administered during a single visit) may be discretionarily limited to avoid significant discomfort to a child.

Note that even without imposing a limit on the number of simultaneous administrations, it may not be possible to construct a schedule in which all the remaining doses can be feasibly scheduled. If a dose for some vaccine cannot be scheduled, the vaccination *series* for that vaccine is considered incomplete.

In certain cases, depending on the child's age and/or the age at which some previous dose is given, it may be beneficial or necessary to prematurely terminate a series. For example, a child normally receives 4 doses of *PCV*, however, the second dose is deemed final if the first dose is administered at age $\geq$ 12 months or the current age is 24-59 months (see Table 1). In either case, the third and fourth doses are unnecessary. This form of *contraindication* can be captured by setting the required gap between the appropriate pair of doses to infinity for the appropriate range for the current age and the age when the earlier dose in the pair is administered. For example, Table 1 demonstrates how one can capture the required spacing and contraindication between the first dose and each subsequent dose of *PCV*.

Given the structure of the rules governing the spacing between doses, it may be possible to cause a contraindication by unnecessarily delaying the administration of some dose. Thus, when constructing a schedule, we are required to maximize the number of completable vaccination series, and among such candidate schedules, maximize the total number of scheduled doses and minimize the total delay from the recommended age of administering these doses.

We next introduce notation for the catch-up scheduling problem:

V	the set of vaccines.
n_v	the total number of doses that constitute the completion of a vaccination series of $v \in V$.
$t_{v,i}^{min}$	the minimum age for administering dose $i \in \{1, \ldots, n_v\}$ of $v \in V$.
$t_{v,i}^{max}$	the maximum age for administering dose $i \in \{1, \ldots, n_v\}$ of $v \in V$.
$t_{v,i}^{rec}$	the recommended age for administering dose $i \in \{1, \ldots, n_v\}$ of $v \in V$.

[1]A "live virus" vaccine is a vaccine that contains a "living" virus that is able to give and produce immunity, usually without causing illness.

$t^{gap}_{v,i,j}(t,t')$ the minimum gap required between dose i and dose $j > i$ of vaccine v when i is administered at age t, and j is being considered for administration at age t'.

$V^{live} \subseteq V$ the set of live vaccines.

t^{live} the minimum gap required between doses of any live vaccines when they are not administered during the same visit.

M the maximum number of simultaneous administrations.

Given a (possibly partial) schedule denoted by s, we use the following notation to define the number and timing of doses scheduled in s:

$\boldsymbol{n}^s_v$ the number of doses of vaccine v that have been scheduled in s.

$\boldsymbol{t}^s_{v,i}$ is the age at which dose $i \in \{1, \ldots, \boldsymbol{n}^s_v\}$ of v is scheduled in s.

$\boldsymbol{t}^s$ an age by which time all doses scheduled in s are administered, i.e., such that $\boldsymbol{t}^s > \boldsymbol{t}^s_{v,i}$ for all $i \in \{1, \ldots, \boldsymbol{n}^s_v\}$ and $v \in V$.

Finally, we introduce some additional notation to define some important characteristics of s:

$m(s,t)$ the number of vaccinations scheduled at age t, i.e., $m(s,t) = \left|\{v \in V : \boldsymbol{t}^s_{v,i} = t \text{ for some } i\}\right|$.

$c(s)$ the number of completable vaccination series, i.e., $c(s) = |\{v \in V : n_v = \boldsymbol{n}^s_v\}|$.

$n(s)$ the total number of doses scheduled, i.e., $n(s) = \sum_{v \in V} \boldsymbol{n}^s_v$.

$d_{v,i}(s)$ the delay from the recommended age of administering dose i of vaccine v, i.e., $d_{v,i}(s) = \max\{0, \boldsymbol{t}^s_{v,i} - t^{rec}_{v,i}\}$.

$d(s)$ the total delay from the recommended age of administering the scheduled doses, i.e., $d(s) = \sum_{v \in V} \sum_{i=1}^{\boldsymbol{n}^s_v} d_{v,i}(s)$.

For a given schedule s, we then define its feasibility and extension as follows.

Definition 1 (Feasibility) *A schedule s is feasible if:*

F1. *s satisfies the time windows for individual doses, i.e. $t^{min}_{v,i} \leq \boldsymbol{t}^s_{v,i} \leq t^{max}_{v,i}$ for all $i \in \{1, \ldots, \boldsymbol{n}^s_v\}$ and $v \in V$,*

F2. *s satisfies the gap requirements between doses of the same vaccine, i.e., $\boldsymbol{t}^s_{v,j} - \boldsymbol{t}^s_{v,i} \geq t^{gap}_{v,i,j}\left(\boldsymbol{t}^s_{v,i}, \boldsymbol{t}^s_{v,j}\right)$ for all $i \in \{1, \ldots, \boldsymbol{n}^s_v\}$, $j \in \{i+1, \ldots, \boldsymbol{n}^s_v\}$ and $v \in V$,*

F3. *s satisfies the gap requirement between doses of live vaccines, i.e.,*

$$\left|\boldsymbol{t}^s_{w,j} - \boldsymbol{t}^s_{v,i}\right| \begin{cases} = 0, & \text{or} \\ \geq t^{live} \end{cases}$$

for all $i \in \{1, \ldots, \boldsymbol{n}^s_v\}$, $j \in \{1, \ldots, \boldsymbol{n}^s_w\}$ and $v, w \in V^{live}$, and

F4. *No more than M doses are scheduled in s at any given age, i.e. $m(s,t) \leq M$ for all t.*

Definition 2 (Extension) *A schedule s' is considered an extension of schedule s, if s' can be obtained from s by scheduling any remaining doses of s on or after $\boldsymbol{t}^s$, i.e.,*

E1. *$\boldsymbol{n}^{s'}_v \geq \boldsymbol{n}^s_v$ for all $v \in V$,*

E2. *$\boldsymbol{t}^{s'}_{v,i} = \boldsymbol{t}^s_{v,i}$ for all $i \in \{1, \ldots, \boldsymbol{n}^s_v\}$ and $v \in V$, and*

E3. $\boldsymbol{t}^{s'}_{v,i} \geq \boldsymbol{t}^s$ *for all* $i \in \{\boldsymbol{n}^s_v + 1, \ldots, \boldsymbol{n}^{s'}_v\}$ *and* $v \in V$.

The catch-up scheduling problem can then be stated as follows: Given a feasible schedule s, the catch-up scheduling problem is one of extending s to a feasible schedule s^* so that

$$\begin{bmatrix} c(s^*) \\ n(s^*) \\ -d(s^*) \end{bmatrix} \geq_L \begin{bmatrix} c(s') \\ n(s') \\ -d(s') \end{bmatrix}$$

for any other feasible extension s' of s. Here we use $\geq_L$ to represent a lexicographical ordering of vectors. Thus, s^* is the best extension of s with respect to (1) the number of completable vaccination series, (2) the number of scheduled doses, and (3) the total delay from the recommended age of administering the scheduled doses, in the stated order of priority.

The catch-up scheduling problem encapsulates many of the complexities of traditional machine scheduling problems. It therefore comes as no surprise that the problem is NP-complete. In fact, it can be shown that the problem remains NP-complete under various simplifications. Despite the similarities to traditional machine scheduling, the catch-up scheduling problem differs since it may not be possible to construct a feasible schedule with all remaining doses and the required separation between doses varies with not only the current age of the child but also, the age at which some previous dose is administered. As a result, one has to employ a multi-level objective that is not typically found in the literature, but at the same time, is ideally suited for DP.

3 Solution Approach

Although the catch-up scheduling problem is NP-complete, the typical size of the problem in practice does not pose a huge challenge. Instead, the challenge we face is in being able to solve these problems consistently within a short amount of time (generally a few seconds).

In this section, we outline a dynamic programming algorithm for solving the catch-up scheduling problem. We start by identifying "dominance" of one schedule over another.

Definition 3 (Dominance) *Given two feasible schedules s_1 and s_2 such that $\boldsymbol{t}^{s_1} = \boldsymbol{t}^{s_2}$, we say s_1 dominates s_2 or $s_1 \succeq s_2$, if we can extend s_1 to a schedule that is at least as good as any schedule obtained from extending s_2.*

Only non-dominated schedules are warranted in the construction of the optimal schedule. Unfortunately, given the complexity of the problem, it is unlikely that there exist efficient necessary conditions for proving dominance of one schedule over another. However, we can find reasonable sufficient conditions by observing that the required spacing between a pair of doses of the same vaccine is generally non-decreasing in the age the first dose in the pair is administered. We first identify within a given schedule s certain "critical" vaccine-dose pairs whose timing in s may prevent some future dose being administered at some age on or after $\boldsymbol{t}^s$:

$$\Psi(s) = \left\{ (v,i) : \begin{array}{c} v \in V, i \in \{1,\ldots,\boldsymbol{n}^s_v\} \text{such that either:} \\ \text{i.} \quad \boldsymbol{t}^s_{v,i} + t^{gap}_{v,i,j}(\boldsymbol{t}^s_{v,i}, t') > t' \text{ for some } j \in \{\boldsymbol{n}^s_v + 1, \ldots, n_v\} \\ \text{and } t' \geq \boldsymbol{t}^s, \text{ or} \\ \text{ii.} \quad \boldsymbol{t}^s_{v,i} + t^{live} > \boldsymbol{t}^s \text{ and } v \in V^{live} \end{array} \right\}.$$

Dominance can then be recognized by using the criteria in the following proposition.

Proposition 1 *If $t^{gap}_{v,i,j}(t,t')$ is non-decreasing in t for all $v \in V$, $i \in \{1,\ldots,n_v\}$, $j \in \{i+1,\ldots,n_v\}$, and t', then for any two feasible schedules s_1 and s_2 such that $\boldsymbol{t}^{s_1} = \boldsymbol{t}^{s_2}$, $s_1 \succeq s_2$ if the following conditions hold:*

D1. $\boldsymbol{n}^{s_1}_v \geq \boldsymbol{n}^{s_2}_v$ *for all* $v \in V$,

D2. $\boldsymbol{n}^{s_1}_v = \boldsymbol{n}^{s_2}_v$ *for all* $v \in V^{live}$ *s.t.* $(v, \boldsymbol{n}^{s_1}_v) \in \Psi(s_1)$,

D3. $\boldsymbol{t}^{s_1}_{v,i} \leq \boldsymbol{t}^{s_2}_{v,i}$ *for all* $(v,i) \in \Psi(s_1)$ *such that* $i \in \{1,\ldots,\boldsymbol{n}^{s_2}_v\}$, *and*

D4. $\sum_{v\in V}\sum_{i=1}^{\boldsymbol{n}^{s_2}_v} d_{v,i}(s_1) \quad \leq \sum_{v\in V}\sum_{i=1}^{\boldsymbol{n}^{s_2}_v} d_{v,i}(s_2)$.

The dominance criteria simply state that s_1 dominates s_2 if (1) s_1 has scheduled at least as many doses as s_2 for each vaccine, (2) s_1 has scheduled the same number of doses as s_2 for any live vaccine whose last dose in s_1 prohibits the scheduling of any other live vaccine at age $\boldsymbol{t}^{s_1}$, (3) the timing of critical doses scheduled in s_1 is no later in s_1 than in s_2, and (4) the total delay in administering doses in common is no worse in s_1 than in s_2.

If the required gap between a pair of doses is non-decreasing in the age the first dose is scheduled, then a schedule with the critical doses scheduled earlier has more flexibility in choosing dates for scheduling future doses. The remaining non-critical doses play no part in determining the timing of any remaining doses but may contribute to the overall quality of the schedule. Given sufficient criteria such as **D1-D4** for efficiently recognizing dominance, we can then use the DP algorithm outlined in Algorithm 1 to construct an optimal extension of a schedule.

For a given schedule s, we define $\tau(s) = \{t_0, t_1, t_2, \ldots\}$ to be the ordered set of time points corresponding to possible ages any remaining dose can be administered starting with $t_0 = \boldsymbol{t}^s$. Starting with s and age t_0, the DP algorithm at iteration k constructs all possible schedules that can be obtained by extending all non-dominated schedules constructed for age t_{k-1} by a single time period. We denote with $\langle s', V', t\rangle$, the new schedule resulting from extending schedule s' by scheduling each vaccine in the set $V' \subseteq V$ at age t and setting $\boldsymbol{t}^{\langle s',V',t\rangle} = t$. Any newly constructed schedule that is infeasible is immediately discarded. Otherwise, $\langle s', V', t\rangle$ is checked for dominance against the candidate set of schedules in $\mathcal{S}(t_k)$ during the **insert** procedure in which any dominated schedules are immediately discarded.

Algorithm 1 The DP Algorithm

Input: schedule s.
Initialize:
$\tau(s) = \{t_0, t_1, t_2, \ldots\}$, $\mathcal{S}_{t_0} \leftarrow \{s\}$, and $\mathcal{S}_{t_k} \leftarrow \{\emptyset\}$ for all $k = 1, \ldots, |\tau(s)|$.
Main Loop:
for $k = 1, \ldots, |\tau(s)|$ **do**
  `/* Iteration` k `*/`
  for all $s' \in \mathcal{S}(t_{k-1})$ **do**
    for all $V' \subseteq V$ s.t. $|V'| \leq M$ **do**
      if $\langle s', V', t_k\rangle$ is feasible **then**
        insert $\langle s', V', t_k\rangle$ into $\mathcal{S}_{t_k}$
Output: **return** schedule $s^* \in \mathcal{S}_{t_{|\tau(s)|}}$ such that $[c(s^*), n(s^*), -d(s^*)] \geq_L [c(s'), n(s'), -d(s')]$ for all $s' \in \mathcal{S}_{t_{|\tau(s)|}}$.

It can be shown by induction that at the start of iteration k of Algorithm 1, $\mathcal{S}_{t_{k-1}}$ contains at least one schedule that can be extended to obtain some best extension of s. Thus, starting with a partial schedule s containing only the past vaccination history of the child and $\boldsymbol{t}^s$ corresponding to the current age of the child, the DP constructs the optimal schedule for administration of the remaining doses. Using the given dominance criteria, we are able to solve most instances within a second and have never encountered any practical instance that took longer than a handful of seconds to solve.

4 A Case Study of Two Scenarios

In this section, we present two solutions obtained for two different real-life scenarios for children requiring catch-up schedules. These cases present varying levels of urgency in terms of how far behind a child has fallen as well as demonstrate the impact of different rules that govern the timing, spacing, and premature termination of a series.

Case 1: A 4 month old child who has received *HepB* at birth and one each of *HepB*, *DTaP*, *Hib*, and *PCV* at 2 months of age.

Case 2: A one year old child without any vaccination.

Figures 1-2 display the different solutions obtained for each of the two scenarios. The first two rows of each chart displays the age and dates for scheduled visits. The first column corresponds to the vaccine line-up. Each box in the chart represents four possible outcomes for a scheduled dose:

AD – an already Administered Dose,

CD – a Catch-up Dose scheduled after the recommended age,

OD – an On-time Dose scheduled during the recommended age, and

PD – a Preemptive Dose scheduled before the recommended age.

At the end of each row we give a tally of doses administered/scheduled out of the total recommended for a vaccination series to be considered completed.

Consider the solution obtained for Case 1 shown in Figure 1. Note also that although this child is 4 months behind for 5 of the 9 vaccines, the schedule has the child catch-up for all but *Rota* by 6 months of age. This is indicated by the trailing **OD** boxes at 6 months of age.

The final solution (Figure 2) displays the solution for Case 2 which is often the standard scenario for internationally adopted or immigrant children presumed not to have received any vaccinations (see (Cohen and Veenstra 2006)). Since the one year old child is assumed not to have received any vaccinations, the standard recommendation would be to vaccinate the child with all 8 vaccines that can be feasibly administered on the current day. However, unless a clinic has many of these in combination, it is unlikely that they would actually administer 8 shots during a single visit. Figure 2 displays the solution when the user chooses to restrict the maximum number of simultaneous administrations to 4.

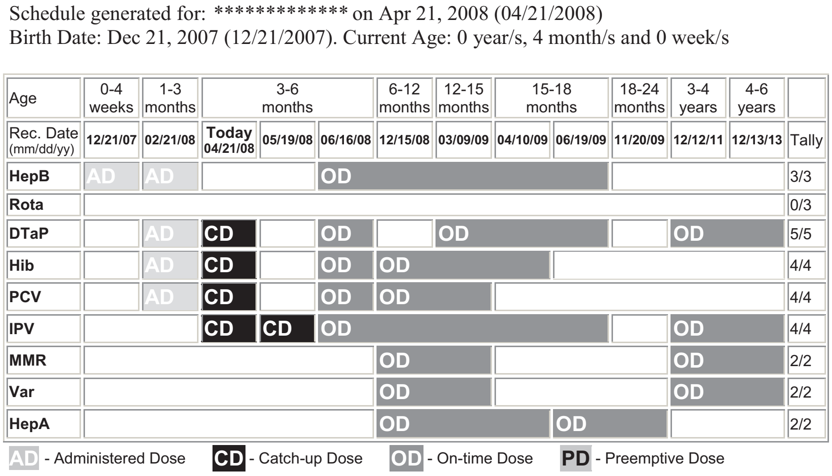

Figure 1: A catch-up schedule constructed for Case 1

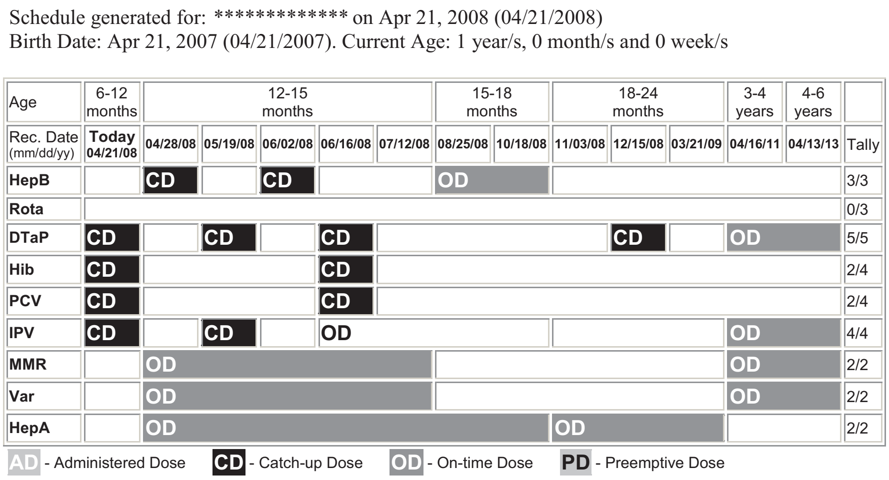

Figure 2: A catch-up schedule constructed for Case 2 when $M = 4$

5 The Scheduler in Practice

The tool was downloaded over 37,110 times from CDC's website during the first of its deployment. The tool has also been referenced in several sections of the latest edition of the AAP Red Book ((Pickering et al. 2009)) considered one of the most definitive guides to immunology and childhood immunization. It has also been featured in over 50 different (online) magazines and news articles including the Washington Post ((Kritz 2008)), U.S. News ((Shute 2008)), Discoveries and Breakthroughs Inside Science ((Ivanhoe Broadcast Network 2008)), AAP News ((Cash 2008)), and was recently featured in CDC's 2009 back-to-school immunization campaign (`www.cdc.gov/Features/CatchUpImmunizations/`).

Several physicians including Dr. Bocchini (chair of the AAP Committee on Infectious Diseases) and Dr. Robert Harrison (Children's Healthcare of Atlanta) who have used the tool have commented that in a busy office they appreciate the rapidity with which decisions can be made by using the tool when a child falls behind in his or her immunizations. They noticed that parents have brought the schedule with them to physician visits and are able to ask appropriate questions and feel they are part of the process. The amount of time saved in determining what vaccines need to be administered when a child is behind and the confidence that the recommended vaccines are administered are major benefits to them. Moreover, physicians feel that the scheduler helps them ensure that children receive vaccines within the recommended guidelines.

6 Conclusions

Healthcare providers are faced on a daily basis with the challenging task of constructing catch-up schedules for childhood immunization. The manual process of constructing such schedules is both difficult and time consuming often resulting in inaccurate or incomplete schedules that can have a detrimental impact on coverage rates and children's health.

In this report, we examine the complicating characteristics of the catch-up scheduling problem and design a DP algorithm that constructs an optimal (with respect to the potential coverage) schedule for a child based on their vaccination history and current age. By observing and exploiting the fact that the required separation between doses of the same vaccine is non-decreasing in the age some previous dose is administered, we derive dominance criteria that are sufficiently tight in practice to solve practically sized problems very quickly.

The tool is currently available for download from CDC's website (`www.cdc.gov/vaccines/scheduler/catchup.htm`) and is being advocated by both the CDC and AAP as a means of encouraging caretakers and providers to take a more proactive role in ensuring timely vaccination coverage of children under their care, and ensuring the accuracy and quality of a catch-up regime. The tool has already provided considerable direction for the rule makers in establishing a more rigorous framework for maintaining consistency in the way current and future rules are stated and in dealing with an infeasible vaccination history. Although it is hard to further quantify the impact of the tool on public policy and/or practice, based on initial feedback and download statistics, we have observed that it has aroused keen interest in the health care community as well as the general public. The fact that the tool is being advocated by both the CDC and AAP additionally testifies to the general acceptance within the healthcare community of the need for such a tool in clinical practice.

Acknowledgments

The author is grateful to Dr. Larry Pickering of from the Centers for Disease Control and Prevention and Prof. Pınar Keskinocak of Georgia Institute of Technology for their vision and support in the development of this tool.

References

Cash, S. 2008. "Online scheduler keeps track of missed immunizations." *AAP News* 29, no. 7 (July).

CDC. 2008. *Recommended schedules and guidelines for catch-up scheduling.* United States: Department of Health and Human Services, National Center for Immunization and Respiratory Diseases. URL: `www.cdc.gov/vaccines/recs/schedules/`.

Cohen, A.L., and D. Veenstra. 2006. "Economic Analysis of Prevaccination Serotesting Compared With Presumptive Immunization for Polio, Diphtheria, and Tetanus in Internationally Adopted and Immigrant Infants." *Pediatrics* 117 (5): 1650–1655.

Cohen, N.J., D.S. Lauderdale, P.B. Shete, J.B. Seal, and R.S. Daum. 2003. "Physician Knowledge of Catch-up Regimens and Contraindications for Childhood Immunizations." *Pediatrics* 111 (May): 925–932.

Holt, E., B. Guyer, N. Hughart, V. Keane, P. Vivier, A. Ross, and D. Strobino. 1996. "The Contribution of Missed Opportunities to Childhood Underimmunization in Baltimore." *Pediatrics*, vol. 97 (April).

Irigoyen, M., S. Findley, P. LaRussa, S. Chen, P. Sternfels, L. Cooper, A. Caesar, and M. Ewing. 2003. "Impact of Immunization Schedule Changes on Missed Opportunities for Vaccination." The 37th National Immunization Conference Chicago, IL.

Ivanhoe Broadcast Network, Inc. 2008. "Keeping Vaccinations On Track." *Discoveries and Breakthroughs Inside Science*. URL: `www.ivanhoe.com/science/story/2008/06/434a.html`.

Kritz, F.L. 2008. "For Parents, an Easier Way to Track Vaccines." *The Washington Post, Health*, no. July 15.

Luman, E.T., M.M. McCauley, S. Stokley, S.Y. Chu, and L.K. Pickering. 2002. "Timeliness of Childhood Immunizations." *American Journal of Preventive Medicine*, vol. 110.

Maciosek, M.V., A.B. Coffield, N.M. Edwards, T.J. Flottemesch, M.J. Goodman, and L.I. Solberg. 2006. "Priorities Among Effective Clinical Preventive Services: Results of a Systematic Review and Analysis." *American Journal of Preventive Medicine* 31:52–61.

Pickering, L.K., C.J. Baker, D.W. Kimberlin, and S.S. Long. 2009. *Red Book: Report of the Committee on Infectious Diseases*. 28th.

Shute, N. 2008. "A New Tool to Manage Your Child's Vaccine Schedule." *U.S. News and World Report, Online*, no. May 27.

Szilagyi, P.G., L.E. Rodewald, S.G. Humiston, R.F. Raubertas, L.A. Cove, C.B. Doane, P.H. Lind, M.S. Tobin, K.L. Roghmann, and C.B. Hall. 1993. "Missed Opportunities for Childhood Vaccination in Office Practices and the Effect on Vacination Status." *Pediatrics*, vol. 91 (January).

Zhou, F., J. Santoli, M.L. Messonnier, H.R. Yusuf, A. Shefer, S.Y. Chu, L. Rodewald, and R. Harpaz. 2005. "Economic Evaluation of the 7-Vaccine Routine Childhood Immunization Schedule in the United States, 2001." *Archives of Pediatrics and Adolescent Medicine* 159:1136–1144.